Recovery Amid Pro-Anorexia: Analysis of Recovery in Social Media

Stevie Chancellor
Georgia Tech
Atlanta, GA 30332
schancellor3@gatech.edu

Tanushree Mitra
Georgia Tech
Atlanta, GA 30332
tmitra3@gatech.edu

Munmun De Choudhury
Georgia Tech
Atlanta, GA 30332
munmund@gatech.edu

ABSTRACT
Online communities can promote illness recovery and improve well-being in the cases of many kinds of illnesses. However, for challenging mental health condition like anorexia, social media harbor both recovery communities as well as those that encourage dangerous behaviors. The effectiveness of such platforms in promoting recovery despite housing both communities is underexplored. Our work begins to fill this gap by developing a statistical framework using survival analysis and situating our results within the cognitive behavioral theory of anorexia. This model identifies content and participation measures that predict the likelihood of recovery. From our dataset of over 68M posts and 10K users that self-identify with anorexia, we find that recovery on Tumblr is protracted - only half of the population is estimated to exhibit signs of recovery after four years. We discuss the effectiveness of social media in improving well-being around anorexia, a unique health challenge, and emergent questions from this line of work.

ACM Classification Keywords
H.4 Information Systems Applications: Miscellaneous

Author Keywords
Tumblr; social media; anorexia; proana; recovery

INTRODUCTION
Participation in social media platforms and online communities is linked to improved well-being and health outcomes [29]. These platforms act as a constantly available and conducive source of information, advice, and support [27, 44, 46]. Community participants also gain a means to learn about the day-to-day aspects of a particular illness, treatment side effects, and managing a personal health challenge [14].

For those struggling with challenging mental health conditions, the role of these platforms in promoting better health is unclear. Examples of conditions include eating disorders, the most prominent of which is anorexia [4]. In their well-established cognitive behavioral theory of anorexia, Fairburn, Shafran and Cooper identify anorexia in individuals with unusual beliefs about food, eating, body shape, weight, and appearance[15]. The Diagnostic and Statistical Manual of Mental Disorders (DSM) identifies anorexia to be the mental illness with the highest mortality rate [1].

On social media platforms, users have reappropriated a general-use space to discuss health topics, and in some cases, these discussions promote dangerous actions. Some individuals that self-identify with anorexia encourage, maintain, and glorify the disorder as a legitimate lifestyle choice rather than a psychosocial disorder [17, 25]. In some extreme cases, individuals in these communities demonstrate pro self-mutilation and pro-suicide sentiments [18, 33]. Individuals with higher exposure to this content are more likely to accept it as normative behavior [63], and the promotion of these dangerous sentiments could be a major challenge towards recovery [2, 24]. Social media sites like Tumblr and Instagram have been the subject of public debate and controversy for the continued presence of anorexia and self-harm content[1,2] [63].

However, many of these platforms also have thriving "recovery" communities. A recovery community is a group of users that discuss the health challenges of mental disorders, promote treatment options, and serve as support for users who are recovering from mental disorders. Recovery communities exist for anorexia in the same social networks that house communities promoting anorexia [70]. Studying both communities, the tensions that develop between them, and the users that frequent them is of interest to researchers drawn to improving anorexia recovery outcomes. However, research has underexplored recovery attempts amid the "anti-recovery" perspective we see on social media [17].

In the light of the above discussion, our central research question involves *examining the role and efficacy of Tumblr as a platform for sustained recovery from anorexia*. We make the following contributions:

(1) Using a hybrid methodology that integrates text processing and human annotations, we identify users who shared anorexia-related content as well as showed signs of recovery in their social media posts.

(2) We develop a robust statistical model based on survival analysis to estimate recovery over time in our user population.

CHI'16, May 07-12, 2016, San Jose, CA, USA
©2016 ACM. ISBN 978-1-4503-3362-7/16/05 $15.00
DOI: http://dx.doi.org/10.1145/2858036.2858246

[1] http://techcrunch.com/2013/06/20/over-a-year-after-new-content-policies-self-harm-social-media-still-thrives/
[2] http://www.huffingtonpost.com/laurenduca/thinspiration-banned-frominstagram_b_3829155.html

Survival analysis offers "time-to-event" data analysis and is widely adopted in the randomized control trial literature [35].

(3) Using our survival analysis over a large dataset of over 10,000 Tumblr users and over 68 million posts, we identify a number of measures that are likely linked to anorexia recovery—body image concerns, behavior, cognition, and affect—based on the cognitive behavioral theory of anorexia [15]. We observe that (a) the estimated time to begin recovery for half of the population is 45 months, and (b) over a six year trajectory, only 56% of our study body is estimated to show signs of recovery on Tumblr.

Our results indicate that anorexia recovery on Tumblr is protracted; our data call into question the long-standing belief that health communities are universally beneficial at encouraging and sustaining positive health outcomes. An important implication of our work is that a challenging psychosocial disorder like anorexia may require significantly different approaches toward encouraging recovery, given the large and vocal presence of a pro-disorder community on social media. Additionally, our work also helps to better understand the potential trajectories of anorexia recovery from pro-anorexic sentiment by leveraging behavioral traces captured in social media, as well as attributes of participation in these platforms that predict this valuable behavior change. We discuss some of the potential implications of using a platform like Tumblr as a health community and the conflicts that arise out of these clashing communities.

Ethics, Privacy, and Disclosure. This paper used publicly accessible Tumblr data to conduct our analysis. No personally identifiable information was used in this study. Quotes taken from users have been slightly modified to protect the identity of these users. Because we did not interact with our subjects and the data is public, we did not seek institutional review board approval. Our work does not make diagnostic claims. Some quotes in this paper are graphic.

BACKGROUND AND MOTIVATION

Studies on Anorexia and Anorexia Recovery

The cognitive behavioral theory of anorexia by Fairburn, Shafran, and Cooper [15] has been widely adopted in the psychology and clinical research literature to examine attributes of recovery in anorexia patients. In this theory, anorexia is characterized by food restrictions through self-starvation and a fear of gaining weight. It also identifies the typical anorexia patient to be introverted, isolated, withdrawn, or depressed, and someone who tries to live up to the expectations of others. Another notable attribute of those with anorexia is Cognitive impairment associated with distorted information processing and dysfunctional styles of reasoning. In the light of these attributes, numerous definitions of recovery, both clinical and psychological, have been proposed [5, 64, 59]. Typically, recovery onset is characterized by a change in attitude toward body image and ingestion, and improved cognitive functioning, self-efficacy, and social cohesion [47, 52].

Clinicians and behavioral scientists have also investigated temporal patterns of anorexia recovery, and scientists have been especially interested in identifying predictors of recovery [23, 59]. In a notable project, Strober et al. [60] performed a longitudinal study of an anorexic population lasting 10-15 years. Following clinical recovery, nearly 76% of the population recovered within 57-79 months. A surprising finding of this work was the small number of identifiable predictors of recovery, compared to predictors of anorexic behavior. Other work has found that the factors contributing to recovery from anorexia are supportive of non-familial relationships, therapy, and maturation [64]. Psychological research on online discussion of anorexia and its recovery has primarily analyzed self-expression and self-presentation [68, 39, 69]. Wolf et. al. [69] analyzed eating disorder, recovery, and unrelated control blogs for linguistic differences; they found that language attributes were particularly predictive for classifying these communities.

However, there are known limitations to the accuracy of self-reported information in anorexia recovery. Self-reported information is commonly gathered through interviews, questionnaires, and writing tasks, and these metrics are used to assess and evaluate progress of an eating disorder or movement into recovery. The social stigma and controversial nature of anorexia, however, make these traditional methods unreliable because feelings of denial, shame, and embarrassment influence response patterns and encourage participants to minimize their feelings [61, 65]. In other words, participants in these clinical trials are known to self-censor or modify their expression to avoid these negative emotions. Additionally, studies of the anorexia recovery experience are frequently limited to clinical in-patients [64]. Many who suffer from anorexia and eating disorders never receive the help of clinicians or mental health professionals because of the stigma associated with these disorders.

While clinical work is a key component to learning more about anorexia recovery, it may miss an important group of people and emotions we believe are accessible through another source—social media. Moreover, the semi-anonymous nature of the Tumblr platform is likely to encourage candid social exchange and self-disclosure, including less conscious manipulation of text [32]. Analyzing text on these platforms offers researchers a non-intrusive and non-reactive way of identifying factors linked to anorexia recovery.

Online Health Communities

A rich body of work in HCI has also revealed how individuals appropriate online platforms and use them to seek health advice and support in unconventional ways. People afflicted by medical conditions often find support via online health communities [14, 54]. One study suggests that 30% of U.S. Internet users have participated in medical or health-related groups [31]. Besides support, these communities serve a range of purposes that include seeking advice [31], connecting with experts and individuals with similar experiences [14, 58, 22], sharing concerns around treatment options [14], sensemaking [42] and understanding professional diagnoses [53], enabling better management of chronic health conditions [43, 29, 30], and fueling discussions with healthcare providers [14]. In this light, approaches to community building have been proposed [20, 66], and the role of participation in such communities toward promoting ailment recovery and coping has been examined in a number of different domains, such as cancer and diabetes [58, 28, 41]. Our investigation in this paper is motivated from this line of research.

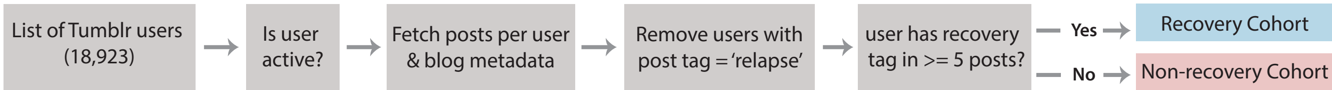

Figure 1: Schematic diagram of our data collection and preparation around identifying the "Recovery" and "Non-recovery" cohorts.

This body of work has not explored recovery and coping for conditions like anorexia. We fill gaps in this space by examining the role and effectiveness of a social media platform harboring communities with both pro-disorder and pro-recovery attitudes in promoting sustained recovery.

Social Media, Health, and Well-Being

Inference of Health Outcomes

Prior research shows the potential to learn about health and well-being through linguistic and behavioral analysis of social media data [50, 45, 12, 26]. De Choudhury et al. [11] analyzed how new mothers' risk to postpartum depression may be detected from Facebook content. MacLean et al. [40] analyzed the text content of a prescription drug abuse forum to learn about recovery and relapse patterns among users. These studies show that valuable information about psychological states is contained within social media content, and that computational techniques can use the linguistic attributes of such content to understand, detect and predict health status. However, the majority of these techniques that rely on supervised learning techniques, such as regression, classification and statistical hypothesis testing, may be inadequate to examine shifts in health states over time. Our work incorporates survival analysis [35] to examine temporal trajectories of the likelihood of change in health status (i.e., recovery from anorexic attitudes in social media).

Social Media and Anorexia

Most prior work in social media and anorexia has used qualitative methods to examine blogs and their discourse about the condition [3, 16, 38, 57]. A few quantitative investigations have recently been conducted [62, 10, 6]; notably, Yom-Tov et al. [70] examined the activity patterns of and interactions between Flickr communities that promote anorexia and those who promote awareness against this condition. In the HCI community, research on eating disorder and anorexia communities is limited. One recent study examined behavior practices of the pro-eating disorder communities before and after Instagram banned tags related to the disorder [6]. None of these works systematically investigated recovery experiences of social media users identifying with anorexia or explored temporal recovery trajectories in social media. Our investigation in this paper is motivated from these lines of research.

DATA

Our investigation uses data from Tumblr, a microblogging service owned by Yahoo!, where users post text and multimedia content to a short-form blog. To build our dataset, we proceeded in three phases.

Phase I: Collecting Pro-Anorexia Data. For our initial data collection, we adopted an approach used in prior work on examination of eating disorders and anorexia on social media sites like Tumblr and Instagram [10]. We first manually examined Tumblr blogs mentioning common eating disorders and their associated anorexia symptomatology tags. Based on a snowball sampling approach during this inspection phase, we obtained an initial list of 28 tags. We examined the co-occurrence of other tags with these seed tags and applied filtering of generic tags (e.g., "fat"). This process expanded our tag list to 304 tags. Examples of these tags include "proana", "anorexia", "thighgap", "thinspiration", "thinspo". We used the Tumblr API to search for posts containing any of these tags. In the process we collected 55,334 *public* English language posts generated by 18,923 unique users.

Phase II: Obtaining Historical Data of Users. In the second phase, we started with our set of 18,923 unique users and retained only those who were still active in Tumblr (since the posts in the first phase were spread over a period of time i.e., between 2008 and 2013, some users were no longer on the platform during the second phase). This left 13,317 active users. For each user we crawled their entire Tumblr history along with their profile information. These include users' activity information (total posts since account creation), the total likes on posts, whether post likes were set to be visible to any other Tumblr user, whether they set 'ask' (questions) to *yes* and 'ask' (questions) by anonymous users to *yes*, and whether their profile is set to share NSFW ("not suitable for work") content. We also gathered post metadata: user id, tags, timestamp, the number of notes (or comments) it received, the number of likes, and users' activity information from their profile (total posts since account creation).

Phase III: Identifying the Recovery Cohort. Our final task involved identifying a candidate set of users who are likely to be recovering from anorexia. We leveraged findings from prior work on expression of recovery tendencies on social media—this literature has identified that social signal of posting to certain specific tags can be a strong indication that a user desires to recover [70, 62, 10]. Such tags have been found to include the identifier/tag "recovery". Based on this observation, we obtained a random sample of 1000 posts containing the regular expression "*recover*". Next, two researchers manually went through the tag list to identify and compile a set of tags co-occurring with the recovery tags for these posts. This was to find only those co-occurring tags which had cues associated with anorexia recovery, e.g., "fighting", "edsoldier". Table 1 lists example tags in our sample. Similarly, we compiled a set of tags with the regular expression "*relaps*" which has been found to be indicative of intent toward anorexia relapse [10]. Taken together, any user who used any of the recovery tags in at least five distinct posts but did not use any of the relapse tags were considered to belong to the "recovery cohort". We refer to the rest of the users as the "non-recovery cohort". Figure 1 illustrates the various steps involved in our data collection and preparation process toward identifying these two cohorts. For illustrative purposes, we enclose example of a (paraphrased) post from a recovery cohort user:

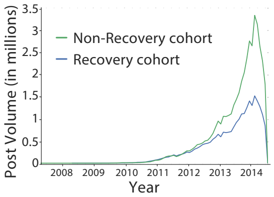

(a) Post Volume by Month

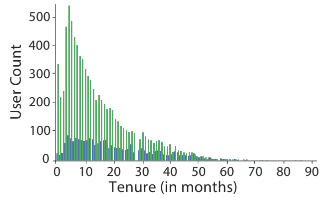

(b) User Tenure by Month

	All	Recovery	Non-Recovery
Total Users	13,317	2,353	10,964
Total Posts	68,380,375	25,710,069	42,670,306
Median Posts / User	1,701.50	4,515	1,276

Figure 2: Summary Statistics describing characteristics of the "Recovery" and the "Non-recovery" cohorts.

> I have a clear mind and a peace I've never known and that's all thanks to recovery. Recovery isn't easy but it is defiantly [sic] worth it. #edrecovery #anarecovery

Similarly, the following is a paraphrased post from a user in our non-recovery cohort:

> I did a half fast yesterday to ease my body into my water fast today, & I've already lost weight! #thinspo #thinsperation

eating disorder recovery	anarecovery	chooserecovery
healthy recovery	pro recovery blog	reasons to recover
recovery fighter	recovery food	recovery intake
recovery record	recovery tips	recoveryisworthit
recoverywarriors	road to recovery	self recovery

Table 1: Sample tags identifying the recovery cohort.

Finally, we evaluated whether the users included in our recovery cohort shared Tumblr content containing signs of recovery. Two researchers familiar with Tumblr and pro-anorexia content online and a clinical psychologist with expertise in eating disorders and self-harm independently evaluated the correctness of the above method. They manually read through the posts of a random sample of 150 users from our recovery cohort, marking (yes/no) whether each user indeed posted content that signaled their desire to recover. Since a single post may not be indicative of the user's inclination to recover, the researchers used a set of three posts per user in the set of 150 users. The raters had high agreement (Cohen's κ=.83) and found that 81% of the users were correctly identified by our method to be in recovery.

Table 2 gives summary statistics of the data in the recovery and non-recovery cohorts. There were 2,353 users in our recovery cohort (25,710,069 posts) while 10,964 in non-recovery (42,670,306 posts). The posts across both cohorts were shared between Feb 20, 2007 and Aug 4, 2014.

METHODS

Models for 'Censored' Data

In human subjects research, one limiting factor across many disciplines is time; study periods are not long enough to observe with certainty whether an event of interest has happened. Clinical research indicates that anorexia recovery is challenging and can take many years [60]. Researchers have tried to analyze the probability of recovery during the study period using conventional statistical techniques (chi-squared test), yet these tests have two problems. First, conventional statistical techniques cannot account for the non-comparability between subjects [21], and second, simply ignoring subjects that never experience the event of study produces biased underestimates of survival [55]. To address this challenge, we borrow techniques from survival analysis methods that have been widely adopted in the randomized control trial literature [35].

Formally, survival analysis is a collection of statistical procedures for analyzing longitudinal data where the outcome of interest is *time until an event occurs* [8]. These techniques are well-suited for scenarios where subjects encounter the event of interest at varying times, cases where they might not even experience the event during the entire observation period, or subjects are lost during the study [49, 67].

In order to employ survival analysis in the context of our research problem, we begin by defining the following terms:

Recovery Event: Our event of interest is the "recovery" event. A user is said to have experienced the "recovery event" if they belong to our identified recovery cohort; that is, they have posted at least five successive posts using a recovery-related tag. We refer to these individuals as *recovering*.

Survival Time: The survival time is the time until an individual experiences (in this case the "recovery event"). We measure this in months. For an individual in the non-recovery cohort who never experience the recovery event, survival time is equal to the entire timespan of the user's collected posts.

Survival Function: The survival function $S(t)$ yields the probability that an individual survives longer than some specified time t, called the *survival probability*. Since events are assumed to occur independently of one another, the probabilities of surviving from one time interval to the next can be multiplied to yield the *cumulative survival probability*.

Hazard Function: The hazard function $H(t)$ gives the probability that an individual who is under observation at a time t will experience the event of interest at that time (known as the *hazard rate*). While the survivor function focuses on cumulative event non-occurrence, the hazard function focuses on the event occurrence and relates to the current event rate [8].

Censoring: Censoring is an important concept in survival analysis and is used to represent cases of missing data. It occurs when we do not know the exact survival time of a user, or information about their survival time is incomplete. For example, consider a Tumblr user with posts until the day our data collection ended (Aug 4, 2014). If the user did not recover by the end of the study, we know that their survival time is at least as long as the study period. However, it is possible that they experience recovery after the study ends. In this scenario, we do not know the exact survival time of the user. Figure 3 illustrates the possible cases where censoring will occur in our study. Here we only have cases where the ex-

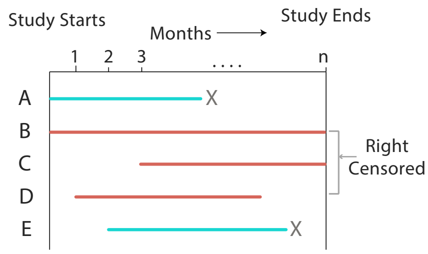

Figure 3: Illustration of censoring (adapted from [36]). The left end of each line corresponds to the timestamp of the first post of the user, while the right is the timestamp of the user's last post in our data. X denotes the recovery event. Users A and E have experienced the "recovery" event, while B, C and D did not. B, C and D's survival time is incomplete at the right side of the study period. These are right-censored data. For B, C and D, we evaluate the censor value = 1, while for A and E, the censor = 0.

act survival time becomes incomplete at the right side of the observation period. This is referred to as *right censored* data.

Why not use linear regression to model the survival time (time taken to experience the "recovery" event, in our case) as a function of a set of predictor variables? First, survival times are typically positive numbers; ordinary linear regression may not be the best choice unless these times are first transformed in a way that removes this restriction. Second, and more importantly, ordinary linear regression cannot effectively handle the censoring of observations. Unlike ordinary regression models, survival methods correctly incorporate information from both censored and uncensored observations in estimating important model parameters.

Statistical Technique

To determine the rate of recovery, we used the Kaplan-Meir estimator for survival analysis [34]. This method provides an estimation of the survival function when the underlying data is censored (as in our case). It estimates the probability of not having the recovery event (i.e. to survive or be in the anorexia state) as a function of time. This is the same as finding the chronological sequence distribution of survival probabilities. The corresponding probability plot is called the survival curve while the tabular representation is referred to as the life table. The median survival time is the time at which one half of the entire cohort recovers.

To find the effect of various factors on the time to recovery, we used the Cox proportional hazards regression model [9]. It is a survival analysis regression method that describes the relationship between the event of interest (in our case, the "recovery event") and the factors that affect the time to that event occurrence. It allows us to estimate the change in the survival probabilities with change in these potential factors. It especially fits our research because Cox modeling does not assume the survival times to follow any particular statistical distribution, unlike most other statistical models.

Measures

We offer a number of content and participation measures for the Cox proportional hazards regression model that predict anorexia recovery. We base these measures on observations in prior literature as well as attributes of the Tumblr platform. Our measures are derived from psychological studies of language use [7] that indicate how different linguistic constructs capture diagnostic information about a wide range of psychological phenomena, ranging from psychiatric disorders, suicidal ideation to responses to a trauma-related upheaval [48, 56]. To characterize content and linguistic constructs of Tumblr content in a systematic and semantically interpretable way, we use the popular psycholinguistic lexicon LIWC (http://www.liwc.net) [51]. Finally, we frame the choice of our content and participation measures in the light of the cognitive behavioral theory of anorexia nervosa described above [15]. Recall that this theory discusses cognitive attributes and thought patterns associated with anorexia nervosa. We propose four measures:

Body Image Concerns: To capture attributes relating to idealized perceptions of body image and the desire for thinness among anorexia sufferers [15, 18], we include measures of the volume of *ingestion*, *body*, and *health* words in the content shared by users — these words and their corresponding categories were obtained from the LIWC dictionary.

Behavior: The cognitive behavioral theory of anorexia finds particular behavioral signatures in anorexic individuals [15]. We consider the following measures of social engagement and nature of interactions that have been examined in prior social media research in the context of mental health [11, 45]. Psychology literature further indicates that these attributes reflect an individual's well-being status [7]. (1) ratio of text to photo posts (prior literature indicates pro-anorexia to have a strong visual expression component [70]); (2) ratio of text to video posts; (3) total number of posts of a user since their account creation; (4) posting entropy, or variability of text versus image versus video post, of a user over time; and (5) length of posts. We further include the following features indicating community-centric interactions: (5) whether a user has 'likes' on their profile posts publicly visible; (6) whether user's profile allows NSFW ("not suitable for work") content sharing; (7) whether the user allows question asking by another Tumblr user; (8) whether they allow questions from anonymous users; and finally (9) whether the users share 'like' information displayed on their profiles.

Cognition: To measure cognitive processes related to anorexia, we again use LIWC: (1) Perception and Regulation: comprising *cognitive mech*, *inhibition*, *insight*, *death*. (2) Temporal References: consisting of the three tense verbs *past*, *present*, and *future tense*. (3) Interpersonal Awareness and Focus: comprising words that are *1st person singular and plural pronoun*, *2nd person pronoun*, and *3rd person* pronouns. (4) Verbal Fluency:the average length of a post. (5) Abusive Language Use:words in LIWC's "Swear" category.

Affect: Finally, we measure affect to characterize emotional expression in Tumblr content. We represent affect as normalized positive affect (PA), computed as the ratio of LIWC words in the positive emotion category, to those in the *negative emotion*, *anger*, *anxiety*, *sadness* categories together.

RESULTS

Survival Analysis

Figure 4 graphs the cumulative probability of experiencing recovery as a function of time. Using the Kaplan Meir estimator for survival analysis, the median time to recovery is

45.6 months. In other words, after 45.6 months, 50% of the user population have not recovered. Probabilities of recovery are also listed for periodic intervals up to 6 years in the life table in Table 2. We see that the probability of recovery 2 years after being on Tumblr is only 16%, while at year 6 it is at 56%. Table 2 underscores that the time course of recovery over the first several years is protracted (i.e., significant lengthening of the time to show signs of recovery in Tumblr content). In fact, as indicated in the survival curve, the likelihood of recovery beyond the 5 year mark (~60 months) is very low – the graph almost shows a flat trend.

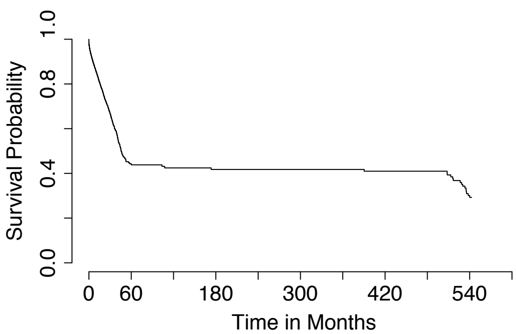

Figure 4: Survival curve showing likelihood of experience of the "recovery" event in our Tumblr data sample of pro-recovery users.

Survival curves can estimate the likelihood that a Tumblr user who has not recovered at a specific time point will remain anorexic for an additional length of time (calculated by dividing the probability of not recovering at Time t_j by the probability of the same at Time t_i, where $j > i$). For example, the probability that a user who did not discuss recovery on Tumblr by Year 2 would remain anorexic for another 2 years is 0.28/0.52 = 53.8%. If she does not recover in 3 years, the probability of remaining anorexic for another 3 years is 0.39/0.56 = 69.6%. As the time of not recovering increases, the likelihood of ever experiencing recovery decreases. This finding aligns with prior work in the anorexia literature where symptomatic recovery patterns had been examined [52].

Time (Years)	Time (Months)	Survival Prob.	Cumulative Prob.	Std. Error
0	0.00	1.00	0.00	0.0003
1	12.00	0.84	0.16	0.0047
2	24.00	0.72	0.28	0.0083
3	36.03	0.61	0.39	0.0141
3.75	45.22	0.51	0.49	0.0235
4	48.01	0.48	0.52	0.0277
5	60.01	0.44	0.56	0.0364
6	72.44	0.44	0.56	0.0382

Table 2: Cumulative probability of remaining in the non-recovery cohort. Median time to recovery is 45.6 months (shown as the shaded row). This is the time when 50% of the users are still expected to not gave recovered.

Recovery Models and Goodness of Fit

In this subsection, we report the goodness of fit of a number of different Cox regression models that estimate the survival probability of users in our data (in other words, the likelihood of experiencing recovery). With the four different categories of measures identified in the Measures section, we report on five models. The first four models correspond to the four measure categories, and the fifth includes all measures from all categories. We refer to them as: `BodyImage`, `Behavior`, `Cognition`, `Affect`, and `Full` models in the rest of this paper. Examining each model separately allows us to compare the four different categories and their role in inferring the likelihood of recovery in our data.

First, we evaluate the goodness of fits of all five of our Cox regression models with *deviance*. Deviance is a measure of the lack of fit to data—lower values are better. Compared to the `Null` model, our models provide considerable explanatory power with significant improvements in deviances. The difference between the deviance of the `Null` model and the deviances of the other models approximately follows a χ^2 distribution with degrees of freedom (df) equal to the number of additional variables in the more comprehensive model. As an example, comparing the deviance of the `Behavior` model with that of the `Null` model, we see that the information provided by the corresponding variables has significant explanatory power: $\chi^2(10, N = 13,317) = 5294.57 - 739.55 = 4,555.02, p < 10^{-6}$. This comparison with the `Null` model is statistically significant after the Bonferroni correction for multiple testing ($\alpha = \frac{0.05}{5}$ as we consider five models). We observe similar deviance results for the `BodyImage`, `Cognition`, `Affect` and `Full` models, with the last model possessing the best fit and highest explanatory power ($\chi^2(26, N = 13,317) = 5294.57 - 214.15 = 5080.42, p < 10^{-10}$).

Table 4 shows the overall Cox model fit by listing the likelihood ratio, Wald and chi-square statistics, and the concordance measure. The `Full` model showed the lowest deviance (refer Table 3), so we report expanded statistics on this model. The tests that generate these statistics are equivalent to the omnibus null hypothesis that all β coefficients are zero. Because the tests shown in Table 4 are statistically significant (Wald statistic $z = 567.4, p < 10^{-15}$) we reject the null hypothesis, indicating that the variables we consider in our `Full` model contribute towards significant explanatory power of estimating the likelihood of recovery. We also note that concordance is a generalization of the area under the receiver operating characteristic (ROC) curve and measures how well a model discriminates between different responses. Specifically, it is the fraction of the pairs of observations in the data, where the observation with the higher survival time has the higher probability of survival predicted by the model [21]. A concordance of greater than 0.5 generally indicates a good prediction ability (the value of 0.5 denotes

Model	Deviance	df	χ^2	p-value
`Null`	5294.57	0		
`BodyImage`	328.63	3	4,965.94	$< 10^{-10}$
`Behavior`	739.55	10	4,555.02	$< 10^{-6}$
`Cognition`	424.92	12	4,869.65	$< 10^{-9}$
`Affect`	682.86	1	4611.71	$< 10^{-6}$
`Full`	214.15	26	5080.42	$< 10^{-10}$

Table 3: Summary of different model fits. Null is the intercept-only model. All comparisons with the Null models are statistically significant after Bonferroni correction for multiple testing ($\alpha = \frac{0.05}{5}$).

		df	p-value
Likelihood ratio test	511.6	26	$< 10^{-15}$
Wald test	567.4	26	$< 10^{-15}$
Score (logrank) test	526.6	26	$< 10^{-15}$
Concordance	0.658	(Std. Err. = 0.007)	
R^2	0.046	(max possible= 0.974)	

Table 4: Summary of fit of the `Full` Cox regression model that estimates likelihood of experience of the recovery event based on our content and activity measures.

no predictive ability). The concordance for the *Full* model is very high (65.8%), and is higher by 14-26% compared to our other models.

What Predicts Recovery?

Predictors of recovery were estimated using the Cox proportional hazards model. Table 5 lists the 26 predictors and their associated β coefficients along with their hazard ratios for the `Full` model. The hazard ratio for a variable denotes the likelihood of a user recovering with one unit increase in the value of the corresponding variable. A hazard ratio smaller than 1 indicates a decreased daily chance of recovering (i.e., increased survival rate), while a hazard ratio larger than 1 indicates an increased chance to recover.

In this section, we discuss how different measures from the four different classes—body image concerns, behavior, cognition, and affect, relate to predictions of recovery.

Body Image. *Reduced body image concerns and increased focus on health and ingestion positively impact the likelihood of anorexia recovery.*

Increased use of *health* and *ingestion* words in Tumblr posts is linked to greater likelihood of recovery. The strongest predictor for recovery in our cohort is discussions about health. Users who mention *health* words increase their monthly hazard of recovery by a factor of 4.1×10^5 ($\beta = 12.9, p < 10^{-3}$). These users express awareness about the severity of anorexia, low body weight, and food restriction as a health risk. Recovering users thus largely use Tumblr as a way to publicly demonstrate their altered perceptions about self-starvation and positive attitudes towards resuming a healthy lifestyle in addition to exhibiting evidence of self-acceptance of their physical appearance:

> “Your weight is only *healthy* when you are.” (recovered)

> “A while ago, a girl asked her friend why she smoked. When she answered ‘i can become anorexic or keep smoking’ and then laughed I just wanted to scream and make her realise that ANOREXIA IS REALLY THAT BITCH.” (recovered)

Other users acknowledge the danger of anorexia and compared the danger of such illnesses to other conditions, like depression or diabetes.

> “I do not encourage any type of depression, eating disorder, bipolar disorder, this is just a blog to represent the mess in my brain. I AM NOT PRO-ANA OR PRO-MIA because this is the stupidest thing I ever heard, eating disorders are ILLNESSES, as well as Cancer, you wouldn’t be PRO-CANCER would you? This is the same..” (recovered)

Improved recovery was also found alongside discussions of ingesting or eating food. Words in this category increased our monthly hazard rate by a factor of 3.7×10^3 ($\beta = 8.2, p = 10^{-3}$). Posts with these words often show a preoccupation with food, nourishment, and eating. It may be surprising that preoccupation with food shows inclination to recovery. Often, however, individuals with anorexia are more focused on their bodies and a focus shift from their physical appearance to food shows movement into recovery [5].

> “Meal 1 at three pm Taco bell bean and cheese caloopa Meal 2 at midnight pasta Additional *foods* Half a chicken quesidia , chips and cheese , donuts Today i finally weighed in at 105 pounds!!! Only ten more to go till my goal of 115 . So proud of myself for coming this far .” (recovered)

> Yes, you should really *feed* yourself. No more of this sitting around for an hour with your cup of *coffee* trying to figure what you want. *Eating* is essential. There is never a time when skipping a *meal* is the right answer. Go *eat*, because you have a life to live. (recovered)

However, posts with *body* words that capture general “body talk” negatively correlate with recovery in our data ($\beta = -19.9, \exp(\beta) = 2.08 \times 10^{-9} (p < 10^{-3})$). These posts offer dietary strategies, fasting regimes, sharing ‘safe’ food or low-calorie foods that users have tried, ways to purge, diet pills, or diuretic abuse. Discussions of specific body parts or how the body looks can also be found here.

> I am on the scale and all what I can think of is how *fat* and ridiculous I am. So I don’t eat, and if I do, I’d purge. But I don’t really care [...] I want to be *skinny* and I don’t care if I die doing this. But at least I will be *skinny* in heaven if not on earth (did not recover)

> You do it because you think your *body* is cute for when your boyfriend picks you up, he feels the *thigh* gap not the *flesh*. But I do it because the voices in my *head* tell me I dont have any other choice. (did not recover)

Behavior. *Greater social engagement, including increased self-disclosure, increases the likelihood of anorexia recovery.*

Next, we observe increased chance of recovery with sharing more text posts than image posts (hazard ratio is increased by one unit; $\beta = 3.5 \times 10^{-3}, p < 10^{-3}$). Longer length of content is linked to increased self-disclosure, which may have positive therapeutic implications [32]. Further, higher post entropy and greater posting activity (in terms of post counts) increases likelihood of recovery; if there is more variability in how users engage with or share content on Tumblr, the chances of recovery increases. This finding aligns with prior literature on social media use and experience of depression [12].

Cognition. *Cognitive mechanisms are a significant predictor in estimating recovery — improved cognitive functioning, ruminative attitude, reduced focus on self-injury and death and lowered self-attentional focus increase the likelihood of anorexia recovery.*

Research has shown that anorexic individuals in recovery have improved cognitive functioning. They speak more insightfully, use more cognitive mechanisms, and inhibition words [39]. Posts that show this self-reflective shift and connection to their own cognition—through use of *cognitive mech*, *inhibition*, and *insight* words—are associated with improved odds of recovery on Tumblr.

> I now *know* I no longer feel fat. Now I no longer *feel incomplete* and *abandoned*. (recovered)

	β	se(β)	z	p value		HR = exp(β)	lower 0.95	upper 0.95
Body Image Concerns								
health	12.92	2.22E+00	5.823	5.77E-09	***	4.08E+05	5.27E+03	3.15E+07
ingestion	8.21	1.95E+00	4.215	2.50E-05	***	3.69E+03	8.10E+01	1.68E+05
body	-19.99	3.89E+00	-5.136	2.80E-07	***	2.08E-09	1.01E-12	4.27E-06
Behavior: Social Engagement								
post entropy	0.11	2.47E-02	4.376	1.21E-05	***	1.11E+00	1.06E+00	1.17E+00
text to photo posts	3.47E-03	8.14E-04	4.258	2.06E-05	***	1.00E+00	1.00E+00	1.01E+00
post length	1.39E-03	2.09E-04	6.645	3.04E-11	***	1.00E+00	1.00E+00	1.00E+00
text to video posts	9.32E-05	9.17E-05	1.017	0.309344		1.00E+00	1.00E+00	1.00E+00
total post count	4.81E-06	1.43E-06	3.352	0.000802	***	1.00E+00	1.00E+00	1.00E+00
Behavior: Community Interactions								
profile allows NSFW content sharing	0.19	2.51E-01	0.743	0.457502		1.21E+00	7.37E-01	1.97E+00
profile allows anon. question asking	0.09	6.89E-02	1.349	0.177429		1.10E+00	9.59E-01	1.26E+00
profile allows question asking	0.06	8.99E-02	0.627	0.530837		1.06E+00	8.87E-01	1.26E+00
profile 'likes' visible	0.00	1.56E-06	0.902	0.367024		1.00E+00	1.00E+00	1.00E+00
profile shares 'like' information	-0.04	5.15E-02	-0.746	0.455379		9.62E-01	8.70E-01	1.06E+00
Cognition: Perception and Regulation								
insight	4.96	2.62E+00	1.892	0.058467	.	1.43E+02	8.37E-01	2.44E+04
cogitive mech	4.87	8.45E-01	5.772	7.82E-09	***	1.31E+02	2.50E+01	6.85E+02
inhibition	3.74	5.87E+00	0.637	0.52437		4.20E+01	4.24E-04	4.16E+06
death	-7.17	9.08E+00	-0.789	0.430146		7.72E-04	1.43E-11	4.17E+04
Cognition: Temporal References								
past tense	11.21	2.10E+00	5.348	8.87E-08	***	7.37E+04	1.21E+03	4.48E+06
present tense	3.57	1.37E+00	2.611	0.009036	**	3.54E+01	2.43E+00	5.15E+02
future tense	1.64	4.00E+00	0.41	0.681981		5.16E+00	2.02E-03	1.32E+04
Cognition: Interpersonal Awareness and Focus								
1st pronoun singular	-8.01	1.74E+00	-4.602	4.18E-06	***	3.33E-04	1.10E-05	1.01E-02
2nd person pronoun	1.11	2.15E+00	0.514	0.607153		3.02E+00	4.46E-02	2.05E+02
1st pronoun plural	-14.31	6.55E+00	-2.184	0.028983	*	6.11E-07	1.62E-12	2.31E-01
3rd person plural	-11.79	3.44E+00	-3.431	0.000601	***	7.62E-06	9.08E-09	6.39E-03
Cognition: Abusive Language Use								
swear	-21.69	6.85E+00	-3.168	0.001533	**	3.81E-10	5.69E-16	2.56E-04
Affect								
Normalized PA	2.51	1.26E+00	2	0.045514	*	1.23E+01	1.05E+00	1.44E+02

Table 5: List of 26 predictors in the `Full` cox regression model. 17 of them are statistically significant (p-values are computed at $\alpha = 0.05$ level followed by Bonferroni correction ($\alpha/26$)). The predictors are listed along with their β weights. $se(\beta)$ is the standard error of the coefficient. z is the Wald statistic which determines whether the corresponding β is statistically significant. The exponentiated β is the hazard ratio (HR). For example, keeping all the predictor variables constant, an additional proportion of *present* words used increases the monthly hazard of recovery by a factor of $e^{\beta} = 35.38$, that is by 3.57 percent.

> You've made a decision. You won't stop, you won't *abandon*. The pain is necessary, especially the pain of hunger. It reassures you that you are STRONG, can withstand anything, that you do NOT have to *admit* to your body, that you don't have to give into its whining. (recovered)

However, a cognition measure that has negative correlation with recovery is discussions of *death*, suicide, or other morbid thoughts ($\beta = -7.17; \exp(\beta) = 7.72 \times 10^{-4}$). Further posts of these users also contain expressions of loneliness, distress, and self-hatred, as well as feelings of social isolation:

> I hate when people say. That teenagers who self harm and think of *suicide* aren't trying to kill themselves but actually trying to live. It might be true for some but not everyone. Because when I was up at night trying to cut open my wrists, hell no I wasn't trying to live.. I was really trying to *die*. (did not recover)

> When I was younger when I first heard about *suicide* I was shocked I didn't know why anybody would *kill* themselves on purpose. all I can saw now is it's funny how fast things can change... (did not recover)

Another strong predictor of recovery is discussions of past events and personal reflections on previous events. Users who discuss past events (indicated by the use of *past* tense words) increase their monthly hazard by a factor of 7.37×10^4 ($\beta = 11.21, p < 10^{-3}$).

> "I feel so bad. I feel guilty for not recovering faster. I want to restrict, binge, and purge, I really want to be thin...I feel guilty for not needing a tube, I've never *had* fortisip[a calorie-dense nutritional beverage often given to patients with eating disorders]...I *was* never put in the wheelchair, I've never passed out, I've never been sick. I don't deserve recovery." (recovered)

Some posts are reflective and can be both regretful or nostalgic of past events, whereas others can be diary-style and may recall a certain period of events. Regardless of topic, using the *past* tense has a link to recovery.

> "have so much i still need to do. 2010 shouldn't be over, after all the stuff I planned, I followed through with nothing. [...]. Next goal: lose 20 pounds before the season starts (putting things in writing gives it meaning, just like telling someone, etc.) and my life ends because no teams will ever want me. Thank God today is over. Night." (recovered)

Heightened use of *future* tense words is linked to increased likelihood of recovery (1.64% increased monthly hazard rating). This suggests that these users are goal-oriented and look forward to the recovery process in the near future [23]. Likelihood of recovery also decreased when also found alongside higher self-attentional focus (indicated by greater use of *first person singular pronouns*), and heightened attention to people and objects (suggested by higher use of third person pronouns). Finally, greater use of abusive language (words in the *swear* category of LIWC) is associated with reduced likelihood of recovery.

Affect. *Increased positive affect and attitude increases the likelihood of recovery.*

Finally, likelihood of recovery is higher in users whose content exhibits a more pronounced hedonic focus on positive emotions and an objective outlook towards life, as indicated by the measure of normalized positive affect (PA) (hazard ratio $\exp(\beta)$ increases by a factor of 1.23, with $\beta = 2.5$, $p < .02$).

> I *love* a life full of eating out with my *loved* ones, sleeping in if that's what I want, getting married and having a family and I can't do these with ana. Let it go away and live your life to the *fullest* and *happiest*. (recovered)

Summary of Findings. Our results show that only half of the cohort we study is estimated to experience the recovery event in ~4 years. Even at the end of a six year period, only 56% of this population is expected to move into recovery. Our Cox regression model reveals several markers predictive of recovery—increased activity on Tumblr, focus on food and nourishment, higher positive affect, and cognitive mechanisms and self-reflection.

DISCUSSION

Our findings show that social media can inform the study of anorexia recovery with information difficult to access through clinical or psychological studies. Using survival analysis, we can identify prospective factors associated with the likelihood of recovery in a large sample; thus we go beyond retrospective analyses of risk factors obtained from self-identified clinical in-patients. Finally, we can longitudinally forecast changes in whether our users will begin to recover.

Implications and Future Directions

Health Practice and Research

Clinical studies of anorexia recovery have often been known to struggle with limited sample of those who suffer from anorexia and may therefore provide biased results in favor of those who actively seek recovery [64]. Our work makes a first attempt at subverting these challenges by focusing on those not traditionally reached through clinical means—self-identified pro-anorexics, or those those who may not seek professional help directly. In general, we believe the observations from our approach can complement clinical work in understanding long-term patterns of anorexia recovery.

We also find text cues that indicate both increased likelihood of moving into recovery as well as markers that indicate maintenance of an eating disorder. Many of these indicators are indeed found in the cognitive behavioral theory of anorexia [15]. We also offer to identify several new markers based on social media community participation and engagement. Our work may lead to future research in behavioral health examining the presence of such cues across other social media. It can also enable clinical researchers examine how the cognitive behavioral theory of anorexia translates to those caught in the conflict between pro-anorexia and anorexia recovery in an online platform.

With new tools and knowledge may also come increased tensions with caregivers and physicians working with patients who use these platforms while seeking assistance with anorexia recovery. With patient consent, physicians could use online social media data like the kind we study here to assist in treatment and management of anorexia of their patients. However, access to this information may have other consequences. Analyzing social media data is an additional workload on already busy caregivers and physicians. Second, if a physician discovers that their patient has expressed self-destructive thoughts on social media, they might be bound by the *duty to treat*. When should a caregiver intervene if they see their patient moving towards relapse on social media? Incorporating social media data into caregiver treatment strategies will need to balance the benefits of more information with the potential workload and ethical considerations.

Social Computing and HCI Research

An impact of our work for social media research is exploring the nuances in how online health communities provide both support and resistance to the anorexia recovery experience. Anorexia recovery is difficult, protracted, and often times accompanied by significant relapses [4]. We see similar difficulties in maintaining recovery in our sample—56% of our cohort remained in recovery through our study.

Our research also shows that not all health communities may necessarily promote recovery or disease management as observed in the prior literature [14]. What could be some of the differences that leads to this outcome? The ecosystem of an otherwise general purpose social media Tumblr and the presence of the anorexia community brings to light the unique situation that these individuals face. In other illness support communities, few people, if any, come to promote the spread or maintenance of a condition like cancer, diabetes, or anxiety. In the context of anorexia, however, pro-anorexia communities exist and coexist in the same spaces as those interested in recovery. A tag as narrow as "anorexia" can contain posts ranging from graphic descriptions of self-harm to extreme diet ideas to positive affirmations and support.

This is not to say, however, that because this population fared no better than clinical results and there are barriers to perfect adoption that social media cannot assist in anorexia recovery. In fact, there are several contexts in which platforms like Tumblr may be facilitating recovery. One way is that the community provides emotional support in times of need and isolation. Often times, those with anorexia suffer alone because of feelings of shame and loneliness [61, 65]. Social media platforms may be a first step to recovery by promoting a sense of togetherness and a place to openly share experiences and emotions with others. Tumblr also provides an emotional "safety valve" [13]. Rather than taking more dangerous or drastic actions (e.g., self-injurious behavior), those suffering can talk through their problems and avoid harmful behavior; social media like Tumblr offer a promising way to engage in disinhibiting and self-disclosing discourse.

Nevertheless, anorexia and associated disorders are unique because body perception and self-esteem are negatively impacted by social comparison and by consumption of images of idealized physical appearance [15]. Social media sites provide affordances around the discoverability of such content that impact one's psychological state [63]. Therefore, in contrast to other health communities that offer support, there are opportunities on social media platforms to make design decisions that not only promote recovery and treatment efforts but also dissipate pro-anorexic behaviors.

(1) Enabling Provisions for Peer Support. Our survival analysis gives probabilistic estimates of the likelihood of recovery in the Tumblr pro-anorexia population. This data can be used by moderators, community managers, and designers to build mechanisms that channel timely support from the greater community. When lowered likelihood of recovery is predicted at a point in the near future, recommendations could appear in the Tumblr dashboard of individuals sharing pro-anorexia content. The system could, for instance, match up recovering anorexics so that they can receive and provide peer motivation. Further, our results demonstrate that positive affect predicts increased likelihood of recovery. Therefore, recovery-related posts from peers with high positive affect and encouraging group solidarity can promoted automatically in these users' social feeds.

(2) Personal Behavior Change Tools. Our work can also assist in building reflective tools that promote behavior change toward and better management of the anorexia recovery experience. These tools include diary-centric tools built on top of social media platforms which would allow an individual to log artifacts (e.g., posts, images, links) they share on Tumblr along their recovery journey. Such aggregated logs, if desired and with consent, could be shared with a caregiver or a clinician, so that the caregivers can track the trajectory of their recovery process and make therapeutic accommodations. Should the posts of an individual signal the presence of a predictor associated with decreased likelihood of recovery, this information could be displayed in the tool's interface, so that the individual can take appropriate action to mitigate the challenge. Candid discourse about health and ingestion correlates with anorexia recovery—the behavior change tools can also leverage these observations in an individual's social feed to encourage expressive writing and self-disclosure as helpful therapeutic practices [37].

Ethical Considerations and Limitations

Ethical Considerations. Anorexia is a complicated and challenging psychosocial disorder, and social media platforms do not have any legal obligation to intervene in pro-anorexia or pro-recovery communities. However, platforms that morally choose to do so must address several ethical considerations when interacting with users. Interventions should be delicately crafted to facilitate positive health outcomes. However, at what point do interventions designed to facilitate recovery become counterproductive or potentially manipulative? It is also important to consider issues arising with algorithmically-based interventions, like surveillance and forecasting of behavior. Privacy also is an important ethical issue with these kind of sensitive data. Although we only gather and analyze public data, social media posts about anorexia are sensitive, and we expect analysis and intervention of this data would be implemented with privacy-preserving measures in mind.

Inferences of Actual Recovery. Although we obtained human ratings on whether a user's content reflected their desire to recover, our findings do not have diagnostic or treatment related claims — we cannot be sure if the individuals actually recovered. Survival analysis gives likelihood values of the experience of the recovery event for the cohort analyzed and does not make individual inferences — as a result, it cannot be used to predict whether a specific individual is going to share content related to recovery. Our methods are to be used at the aggregate level to understand the "well-being" of community and to make subsequent provisions for individuals whose content exhibits cues linked to increased or decreased likelihood of recovery.

Self-Presentation. We also recognize that content shared on Tumblr may be self-censored by users that aligns with their personality traits and perceptions of their social audience on the platform. People may not use the tag "pro-recovery" to report about their desire or experience of doing so and would not be accounted for by our approach. Our method solely depends on the content in the posts, and we cannot account for inter-individual differences due to personality attributes, self-presentation, identity, or self-censorship [19]. However, since we make community-wide claims about recovery, we believe our findings generalize to the broader userbase. Moreover, Tumblr provides a fairly anonymous platform of exchange, hence we believe our sample to be less concerned about self-presentation than a platform like Facebook.

Stages of Recovery. We were interested in the broad notion of sharing about the experience or desire of recovery from anorexia, but did not characterize the nuances of this recovery process. The Transtheoretical Model (TTM) is a method that is often adopted to identify the various stages of recovering from an ailment [40], although in most prior work the stages have been identified through qualitative or manual labeling. In the future we are interested in deploying our survival analysis with a hierarchical or staged approach that can help us reveal these recovery stages and therefore help characterize better how social media is used by individuals to facilitate recovery from anorexia. TTM can also help us identify suitable points in time during the course of recovery where an individual might have the most self-regulation to be receptive to an intervention that can help them move towards recovery.

CONCLUSION

In this paper, we provided a large-scale quantitative analysis of anorexia recovery on social media on Tumblr. We use survival analysis to understand the likelihood of whether ~13K pro-anorexic users on Tumblr would recover. Only half of our cohort shows likelihood of recovery after 45 months, and a vast minority (44%) is not estimated to recover even at the end of six years. Using the cognitive behavioral theory of anorexia, we also identify several linguistic and behavioral factors that may indicate an increased likelihood to recover. Given that we observe recovery outcomes on Tumblr to be less than those observed in existing clinical studies, we discuss how online communities may design provisions to facilitate recovery. Our work is a first step towards understanding vulnerable online communities like anorexia recovery, and we believe that our research helps both heath researchers and social media designers bring support to those in need.

ACKNOWLEDGEMENTS

The authors thank Erica Goodman for her help in rating posts. Chancellor and De Choudhury were partly supported through an NIH grant # 1R01GM11269701. Mitra was funded by a DARPA Award # DARPA-W911NF-12-1-0043.

REFERENCES

1. American Psychiatric Association and others. 2003. *Diagnostic and Statistical Manual of Mental Disorders:: DSM-5*. ManMag.
2. Anna M Bardone-Cone and Kamila M Cass. 2006. Investigating the impact of pro-anorexia websites: A pilot study. *European Eating Disorders Review* 14, 4 (2006), 256–262.
3. Dina LG Borzekowski, Summer Schenk, Jenny L Wilson, and Rebecka Peebles. 2010. e-Ana and e-Mia: A Content Analysis of Pro–Eating Disorder Web Sites. *American journal of public health* 100, 8 (2010), 1526.
4. Hilde Bruch. 1973. *Eating disorders: Obesity, anorexia nervosa, and the person within*. Vol. 5052. Basic Books.
5. Cynthia M Bulik, Nancy D Berkman, Kimberly A Brownley, Jan A Sedway, and Kathleen N Lohr. 2007. Anorexia nervosa treatment: a systematic review of randomized controlled trials. *International Journal of Eating Disorders* 40, 4 (2007), 310–320.
6. Stevie Chancellor, Jessica Pater, Trustin Clear, Eric Gilbert, and Munmun De Choudhury. 2016. #thyghgapp: Instagram Content Moderation and Lexical Variation in Pro-Eating Disorder Communities. In *Proceedings of the 2016 Conference on Computer Supported Cooperative Work (CSCW)*. ACM.
7. Cindy Chung and James W Pennebaker. 2007. The psychological functions of function words. *Social communication* (2007), 343–359.
8. TG Clark, MJ Bradburn, SB Love, and DG Altman. 2003. Survival analysis part I: basic concepts and first analyses. *British journal of cancer* 89, 2 (2003), 232–238.
9. David R Cox. 1992. Regression models and life-tables. In *Breakthroughs in Statistics*. Springer, 527–541.
10. Munmun De Choudhury. 2015. Anorexia on Tumblr: A Characterization Study. In *Proc. Digital Health*.
11. Munmun De Choudhury, Scott Counts, Eric Horvitz, and Aaron Hoff. 2014. Characterizing and Predicting Postpartum Depression from Facebook Data. In *Proc. CSCW*.
12. Munmun De Choudhury, Michael Gamon, Scott Counts, and Eric Horvitz. 2013. Predicting depression via social media. In *Proc. ICWSM*.
13. T Emmens and A Phippen. 2010. Evaluating Online Safety Programs. *Harvard Berkman Center for Internet and Society.[23 July 2011]* (2010).
14. Gunther Eysenbach, John Powell, Marina Englesakis, Carlos Rizo, Anita Stern, and others. 2004. Health related virtual communities and electronic support groups: systematic review of the effects of online peer to peer interactions. *Bmj* 328, 7449 (2004), 1166.
15. Christopher G Fairburn, Roz Shafran, and Zafra Cooper. 1999. A cognitive behavioural theory of anorexia nervosa. *Behaviour Research and Therapy* 37, 1 (1999), 1–13.
16. Rachel A Fleming-May and Laura E Miller. 2010. I'm scared to look. But I'm dying to know: Information seeking and sharing on Pro-Ana weblogs. *Proc. ASIST* 47, 1 (2010), 1–9.
17. Nick Fox, Katie Ward, and Alan O'Rourke. 2005. Pro-anorexia, weight-loss drugs and the internet: an "anti-recovery" explanatory model of anorexia. *Sociology of Health & Illness* 27, 7 (2005), 944–971.
18. Jeff Gavin, Karen Rodham, and Helen Poyer. 2008. The presentation of "pro-anorexia" in online group interactions. *Qualitative Health Research* 18, 3 (2008), 325–333.
19. Erving Goffman. 1959. The presentation of self in everyday life. (1959).
20. Andrea Grimes, Brian M Landry, and Rebecca E Grinter. 2010. Characteristics of shared health reflections in a local community. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*. ACM, 435–444.
21. Frank E Harrell. 2001. *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*. Springer.
22. Andrea Hartzler and Wanda Pratt. 2011. Managing the personal side of health: how patient expertise differs from the expertise of clinicians. *Journal of medical Internet research* 13, 3 (2011).
23. David B Herzog, David J Dorer, Pamela K Keel, Sherrie E Selwyn, Elizabeth R Ekeblad, Andrea T Flores, Dara N Greenwood, Rebecca A Burwell, and Martin B Keller. 1999. Recovery and relapse in anorexia and bulimia nervosa: a 7.5-year follow-up study. *Journal of the American Academy of Child & Adolescent Psychiatry* 38, 7 (1999), 829–837.
24. Sharlene Hesse-Biber, Patricia Leavy, Courtney E Quinn, and Julia Zoino. 2006. The mass marketing of disordered eating and eating disorders: The social psychology of women, thinness and culture. In *Women's studies international forum*, Vol. 29. Elsevier, 208–224.
25. Hans Wijbrand Hoek. 2006. Incidence, prevalence and mortality of anorexia nervosa and other eating disorders. *Current opinion in psychiatry* 19, 4 (2006), 389–394.
26. Christopher M Homan, Naiji Lu, Xin Tu, Megan C Lytle, and Vincent Silenzio. 2014. Social structure and depression in TrevorSpace. In *Proc. CSCW*.
27. Yan Hong, Ninfa C Peña-Purcell, and Marcia G Ory. 2012. Outcomes of online support and resources for cancer survivors: A systematic literature review. *Patient education and counseling* 86, 3 (2012), 288–296.
28. Mette Terp Høybye, Christoffer Johansen, and Tine Tjørnhøj-Thomsen. 2005. Online interaction. Effects of storytelling in an internet breast cancer support group. *Psycho-Oncology* 14, 3 (2005), 211–220.

29. Jina Huh and Mark S Ackerman. 2012. Collaborative help in chronic disease management: supporting individualized problems. In *Proc. CSCW*.

30. Jina Huh, Leslie S Liu, Tina Neogi, Kori Inkpen, and Wanda Pratt. 2014. Health vlogs as social support for chronic illness management. *ACM Transactions on Computer-Human Interaction (TOCHI)* 21, 4 (2014), 23.

31. Grace J Johnson and Paul J Ambrose. 2006. Neo-tribes: The power and potential of online communities in health care. *Commun. ACM* 49, 1 (2006), 107–113.

32. Adam N Joinson. 2001. Self-disclosure in computer-mediated communication: The role of self-awareness and visual anonymity. *European Journal of Social Psychology* 31, 2 (2001), 177–192.

33. Adrienne S Juarascio, Amber Shoaib, and C Alix Timko. 2010. Pro-eating disorder communities on social networking sites: a content analysis. *Eating disorders* 18, 5 (2010), 393–407.

34. Edward L Kaplan and Paul Meier. 1958. Nonparametric estimation from incomplete observations. *Journal of the American statistical association* 53, 282 (1958), 457–481.

35. John P Klein and Melvin L Moeschberger. 2003. *Survival analysis: techniques for censored and truncated data*. Springer Science & Business Media.

36. David G Kleinbaum. 1998. Survival Analysis, a Self-Learning Text. *Biometrical Journal* 40, 1 (1998), 107–108.

37. Stephen J Lepore and Joshua M Smyth. 2002. *The writing cure: How expressive writing promotes health and emotional well-being.* American Psychological Association.

38. Stephen P Lewis and Alexis E Arbuthnott. 2012. Searching for thinspiration: the nature of internet searches for pro-eating disorder websites. *Cyberpsychology, Behavior, and Social Networking* 15, 4 (2012), 200–204.

39. Elizabeth J Lyons, Matthias R Mehl, and James W Pennebaker. 2006. Pro-anorexics and recovering anorexics differ in their linguistic Internet self-presentation. *Journal of Psychosomatic Research* 60, 3 (2006), 253–256.

40. 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. CSCW*.

41. Lena Mamykina, Andrew D Miller, Elizabeth D Mynatt, and Daniel Greenblatt. 2010. Constructing identities through storytelling in diabetes management. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1203–1212.

42. Lena Mamykina, Drashko Nakikj, and Noemie Elhadad. 2015. Collective Sensemaking in Online Health Forums. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 3217–3226.

43. Jennifer Mankoff, Kateryna Kuksenok, Sara Kiesler, Jennifer A Rode, and Kelly Waldman. 2011. Competing online viewpoints and models of chronic illness. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 589–598.

44. Phoenix KH Mo and Neil S Coulson. 2008. Exploring the communication of social support within virtual communities: A content analysis of messages posted to an online HIV/AIDS support group. *Cyberpsychology & behavior* 11, 3 (2008), 371–374.

45. Elizabeth L Murnane and Scott Counts. 2014. Unraveling abstinence and relapse: smoking cessation reflected in social media. In *Proc. CHI*.

46. Mark W Newman, Debra Lauterbach, Sean A Munson, Paul Resnick, and Margaret E Morris. 2011. It's not that I don't have problems, I'm just not putting them on Facebook: challenges and opportunities in using online social networks for health. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*. ACM, 341–350.

47. Greta Noordenbos, Anna Oldenhave, Jennifer Muschter, and Nynke Terpstra. 2002. Characteristics and treatment of patients with chronic eating disorders. *Eating disorders* 10, 1 (2002), 15–29.

48. Thomas E Oxman, Stanley D Rosenberg, and Gary J Tucker. 1982. The language of paranoia. *The American Journal of Psychiatry* (1982).

49. Mahesh KB Parmar and David Machin. 1995. *Survival analysis: a practical approach*. John Wiley & Sons Chichester.

50. Michael J Paul and Mark Dredze. 2011. You Are What You Tweet: Analyzing Twitter for Public Health.. In *Proc. ICWSM*.

51. James W Pennebaker, Tracy J Mayne, and Martha E Francis. 1997. Linguistic predictors of adaptive bereavement. *Journal of personality and social psychology* 72, 4 (1997), 863.

52. Kathleen M Pike. 1998. Long-term course of anorexia nervosa: response, relapse, remission, and recovery. *Clinical psychology review* 18, 4 (1998), 447–475.

53. Jenny Preece and Diane Maloney-Krichmar. 2005. Online communities: Design, theory, and practice. *Journal of Computer-Mediated Communication* 10, 4 (2005), 00–00.

54. Shelly Rodgers and Qimei Chen. 2005. Internet community group participation: Psychosocial benefits for women with breast cancer. *Journal of Computer-Mediated Communication* 10, 4 (2005), 00–00.

55. Philip Rowe. 2007. *Essential statistics for the pharmaceutical sciences*. John Wiley & Sons.

56. Stephanie Rude, Eva-Maria Gortner, and James Pennebaker. 2004. Language use of depressed and depression-vulnerable college students. *Cognition & Emotion* 18, 8 (2004), 1121–1133.

57. Leslie Regan Shade. 2003. Weborexics: The Ethical Issues Surrounding Pro-Ana Websites. *SIGCAS Comput. Soc.* 33, 4 (Dec. 2003).

58. Meredith M Skeels, Kenton T Unruh, Christopher Powell, and Wanda Pratt. 2010. Catalyzing social support for breast cancer patients. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 173–182.

59. Nalini M Srinivasagam, Walter H Kaye, Katherine H Plotnicov, Catherine Greeno, Theodore E Weltzin, and Radhika Rao. 1995. Persistent perfectionism, symmetry, and exactness after long-term recovery from anorexia nervosa. *American Journal of Psychiatry* 152, 11 (1995), 1630–1634.

60. Michael Strober, Roberta Freeman, and Wendy Morrell. 1998. The long-term course of severe anorexia nervosa in adolescents: survival analysis of recovery, relapse, and outcome predictors over 10-15 years in a prospective study. *International Journal of Eating Disorders* 22 (1998), 339–60.

61. Sarah Swan and Bernice Andrews. 2003. The relationship between shame, eating disorders and disclosure in treatment. *British journal of clinical psychology* 42, 4 (2003), 367–378.

62. Shabbir Syed-Abdul, Luis Fernandez-Luque, Wen-Shan Jian, Yu-Chuan Li, Steven Crain, Min-Huei Hsu, Yao-Chin Wang, Dorjsuren Khandregzen, Enkhzaya Chuluunbaatar, Phung Anh Nguyen, and others. 2013. Misleading health-related information promoted through video-based social media: anorexia on YouTube. *Journal of medical Internet research* 15, 2 (2013).

63. Stephanie Tierney. 2006. The dangers and draw of online communication: Pro-anorexia websites and their implications for users, practitioners, and researchers. *Eating Disorders* 14, 3 (2006), 181–190.

64. Federica Tozzi, Patrick F Sullivan, Jennifer L Fear, Jan McKenzie, and Cynthia M Bulik. 2003. Causes and recovery in anorexia nervosa: The patient's perspective. *International Journal of Eating Disorders* 33, 2 (2003), 143–154.

65. Kelly Bemis Vitousek, Jennifer Daly, and Christopher Heiser. 1991. Reconstructing the internal world of the eating-disordered individual: Overcoming denial and distortion in self-report. *International Journal of Eating Disorders* 10, 6 (1991), 647–666.

66. Paul Wicks, Michael Massagli, Jeana Frost, Catherine Brownstein, Sally Okun, Timothy Vaughan, Richard Bradley, and James Heywood. 2010. Sharing health data for better outcomes on PatientsLikeMe. *Journal of medical Internet research* 12, 2 (2010).

67. John B Willett and Judith D Singer. 1993. Investigating onset, cessation, relapse, and recovery: why you should, and how you can, use discrete-time survival analysis to examine event occurrence. *Journal of consulting and clinical psychology* 61, 6 (1993), 952.

68. Markus Wolf, Jan Sedway, Cynthia M Bulik, and Hans Kordy. 2007. Linguistic analyses of natural written language: Unobtrusive assessment of cognitive style in eating disorders. *International journal of eating disorders* 40, 8 (2007), 711–717.

69. Markus Wolf, Florian Theis, and Hans Kordy. 2013. Language Use in Eating Disorder Blogs Psychological Implications of Social Online Activity. *Journal of Language and Social Psychology* 32, 2 (2013), 212–226.

70. Elad Yom-Tov, Luis Fernandez-Luque, Ingmar Weber, and Steven P Crain. 2012. Pro-anorexia and pro-recovery photo sharing: A tale of two warring tribes. *Journal of medical Internet research* 14, 6 (2012).