

Anorexia on Tumblr: A Characterization Study

Munmun De Choudhury
School of Interactive Computing
Georgia Institute of Technology
Atlanta, GA 30308
munmund@gatech.edu

ABSTRACT

Eating disorders, such as anorexia nervosa are a major health concern affecting many young individuals. Given the extensive adoption of social media technologies in the anorexia affected demographic, we study behavioral characteristics of this population focusing on the social media Tumblr. Aligned with observations in prior literature, we find the presence of two prominent anorexia related communities on Tumblr — pro-anorexia and pro-recovery. Empirical analyses on several thousand Tumblr posts show use of the site as a media-rich platform replete with triggering content for enacting anorexia as a lifestyle choice. Through use of common pro-anorexia tags, the prorecovery community however attempts to “permeate” into the pro-anorexia community to educate them of the health risks of anorexia. Further, the communities exhibit distinctive affective, social, cognitive, and linguistic style markers. Compared with recovering anorexics, pro-anorexics express greater negative affect, higher cognitive impairment, and greater feelings of social isolation and selfharm. We also observe that these characteristics may be used in a predictive setting to detect anorexia content with ~80% accuracy. Based on our findings, clinical implications of detecting anorexia related content on social media are discussed.

Categories and Subject Descriptors

H5.3 [Group and Organization Interfaces: Asynchronous interaction]: Web-based interaction

Keywords

anorexia, eating disorder, health, social media, tumblr.

1. INTRODUCTION

There are many adolescents and young adults today who are challenged by eating disorders and its associated behavioral health concerns. Statistics from the National Eating

Disorder Association (NEDA)¹ report that in the United States, more than 20 million women and 10 million men suffer from a clinically significant eating disorder at some time in their life. This includes anorexia nervosa, bulimia nervosa, binge eating disorder, etc. Anorexia nervosa, particularly, is the most common form of eating disorder, and is classified as a mental health disorder (DSM-IV). It is typically characterized by food restrictions in the form of self-starvation and a fear of gaining weight, and often adopted as a legitimate alternative lifestyle choice [18]. In fact, anorexia has the highest mortality rate of any mental illness [10].

However, there has been significant lack of resources to combat eating disorders and their devastating consequences [10]: only 1 in 10 individuals with anorexia receive treatment. Further, a line of research argues that the proliferation and pervasiveness of online technologies over the last decade has led to an increase in the prevalence of the condition [16, 18]. This is because the use of these tools often exposes adolescent girls to thin idealized female images, typically known as “thinspiration” content. Moreover, these platforms have created thriving communities wherein individuals can find and disseminate material to support, endorse, or describe the progression and maintenance of anorexia.

Attempts are being made by online forums to tackle issues of self-harm, and the like: aspects associated with anorexic populations². However most of them only act when infringing accounts are brought to attention by users, or deploy rudimentary keyword detection approaches to ban “triggering” content. Some of them (e.g. YouTube) also choose to censor such content, while others provide a general public service announcement (e.g. Tumblr). Given the growing volume of such content, there is therefore a need to devise more effective techniques to improve these efforts.

The association between internet use and adoption of pro-anorexic lifestyle is well-studied. However, a relatively less explored area involves the examination of the use of social media by the affected demographic. According to a Pew Internet study, 81% of online teens aged 12-17 use some kind of social media [14]. Note that this is also a demographic that is known to be most affected by anorexia [10]. To this end, in this paper we develop computational approaches to identify behavioral characteristics of communities that share anorexia related content.

We focus on Tumblr (<http://www.tumblr.com/>) as the social media platform of our study. Of late, Tumblr has strug-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
DH'15, May 18–20, 2015, Florence, Italy.
Copyright © 2015 ACM 978-1-4503-3492-1/15/05 ...\$15.00.
<http://dx.doi.org/10.1145/2750511.2750515>.

¹<http://www.nationaleatingdisorders.org/>

²<http://staff.tumblr.com/post/18132624829/self-harm-blogs>

gled with policies around self-harm². This context serves as one of the motivations for the choice of this platform for our study. Additionally, the younger userbase of Tumblr (Tumblr’s characteristic audience are teen and college-aged user segments³) can provide new opportunities to understand their involvement in sharing anorexia content on social media. Specifically, Tumblr can let us examine how use of this emergent microblogging platform is distinct from the use of other online resources on anorexia, as explored in prior literature [28, 29].

Specifically, in this paper we observe that anorexia related content on Tumblr surfaces from two prominent communities: (1) The pro-anorexia (popularly known as “proana”) community, which typically comprises blogs that glorify, encourage, or promote anorexia, including self-mutilation and self-destructive thoughts; and (2) the pro-recovery community who use the platform to seek support and share recovery experiences with others, or to alter perceptions advocating anorexia as a lifestyle choice. Since these communities have different goals behind use of social media [3, 18], we believe characterization of anorexia content needs to contrast these communities. We address the following research questions:

RQ 1 What are the characteristics of the proana and pro-recovery communities, in terms of their affective, behavioral, cognitive, and linguistic style attributes?

RQ 2 Are the social media characteristics of these communities predictive of whether a post disseminates proana or pro-recovery information? Generally, can these characteristics be used to detect anorexia related content on social media?

Our findings suggest that the pro-anorexia community uses this microblogging platform Tumblr to share image-rich graphic and “triggering” content around internalization of thin body ideals, as well as for the maintenance of anorexic lifestyles. Interestingly the pro-recovery community, through the use of common pro-anorexic tags, attempts to “permeate” into the pro-anorexia community. We conjecture the pro-recovery community thus serves an important social responsibility by encouraging adoption of recovery therapy, and highlighting the health risks of anorexia.

Further, we find that by leveraging the language used in Tumblr posts of the pro-anorexia and pro-recovery communities, we can measure various affective, social, cognitive, and stylistic elements of their behavior. Compared with recovering anorexics, pro-anorexics express greater negativity, cognitive impairment and self-occupation, and greater feelings of social isolation and self-harm. We also observe that these characteristics may be used in a predictive setting to identify pro-anorexia or pro-recovery content, as well as more generally to distinguish between anorexia related and other content on Tumblr (~80% accuracy). We discuss the implications of our findings in the identification of social media communities vulnerable to anorexia or attempting to recover from anorexia, and in the design of interventions that can bring help to these two distinctive communities. We also present design suggestions for social media based health interventions, that can tailor behavioral therapy for proana and pro-recovery individuals. Finally, we also discuss our social and ethical responsibilities towards pursuing this line of research.

³<http://www.comscore.com/Insights/Blog>

2. BACKGROUND AND RELATED WORK

2.1 Studies on Anorexia

Psychology researchers have long examined the behavioral characteristics of anorexia and eating disorder behaviors [27]. Lyons et al [15] explored the presence of linguistic markers of differences in Internet self-presentation of self-identified pro-anorexics who defended anorexia as a lifestyle, and self-identified anorexics in recovery. On similar lines, Wolf et al. [27] implemented a writing task for eating disorder inpatients and investigated the language they used.

The above research suggests that automated analysis of natural language can offer a novel and reliable tool to study a stigmatic, sensitive, and socially disputed condition like anorexia. However, note that, assessments via self-report (interviews, questionnaires), a method adopted by many of the above prior work, have been found to be inadequate for studying eating disorder affected individuals. This is because denial and minimization due to feelings of shame and embarrassment influence response patterns [24, 26]. Behavior assessed from language shared in social media and spanning several thousand individuals are less subject to conscious manipulation. Therefore it can offer an approach to non-intrusive and non-reactive assessment.

Prior work focusing on social media and sharing of anorexia related content is relatively limited. Close to our work is another work by Wolf et al. [28] wherein the authors analyzed pro-eating disorder, recovery blogs, and control blogs through quantitative text analyses. Yom-Tov et al. [29] examined the activity patterns of and interactions between proana and pro-recovery Flickr communities. Syed-Abdul et al. [25] studied how misleading information was shared through pro-anorexia videos on YouTube.

The findings in this body of literature are valuable, and have helped us frame our research questions. We attempt to extend this line of research by examining how newly emergent social media like Tumblr, that are particularly pervasive in the age group where anorexia is prevalent, are being appropriated for the purpose. Further, we investigate whether social media activity bears distinctive markers that can be used to automatically identify anorexia related content.

2.2 Social Media and Health

Recently, a growing body of work has increasingly been interested in understanding how large-scale social media activities can be used to understand, infer, and improve the wellbeing of people [19], including conditions and symptoms related to diseases [22], substance abuse and distress [8, 17], postpartum depression [7], mental health [21, 11], insomnia [12] and so on. A common thread in this body of research is how computational techniques may be applied to naturalistic data that people share on today’s online social platforms to make sense of their health behaviors. This paper builds on this promising new line of research.

Moreover, a line of research investigates the impact of increasing use of online platforms, such as websites and blogs, in discourse around pro eating disorders [2, 3, 16]. It is known that a segment of population on these platforms take self-destruction to unimaginable extremes [18], with users encouraging the sharing of content that promote negative perceptions of body image [13], or promoting pro-suicide sentiments. We believe that by pursuing the potential of automatically identifying defining characteristics

of pro-anorexic and pro-recovery content shared on social media, we will be able to frame directions on guiding interventions that can help address the health challenges around anorexia.

3. DATA

Tumblr is a microblogging service owned by Yahoo! that allows users to post text and multimedia content to a short-form blog. It is the 6th most popular social networking service⁴. We use the official Tumblr API (<http://www.tumblr.com/docs/en/api/v2>) to collect 55,334 *public* English language posts from 18,923 unique blogs, which we refer to as users. Our data collection proceeded in the following manner: We first manually examined a number of Tumblr blogs mentioning common eating disorder and anorexia symptomatology tags. Based on a snowball sampling approach during this inspection phase we obtained an initial list of 28 tags. Next we performed an initial week long data crawl using these 28 tags. From this crawl, we followed the tags that co-occurred frequently with these 28 tags (in terms of high point wise mutual information, PMI)—this gave us a sample of 304 tags. After filtering out generic tags like “life”, “fat”, “personal”, “food” etc. or tags that do not uniquely characterize anorexia or related disorders, our final crawl utilized a set of 72 tags (ref. Table 1).

anorexia	anoreixa	bulima
eating disorder	anamia	ed
thinspo	bulimic	pro ana
eating disorder recovery	anarecovery	anorexic
thinspiration	promia	thigh gap
eating disorder recovery	mia	anorexique
anabuddyneeded	anadvice	thinspiration tips
proanania	actuallyanorexic	eating disorder tw

Table 1: Sample tags (out of 72) used for Tumblr crawl.

We collected Tumblr posts, the associated tags of the post, title of the post, its timestamp, the number of notes it got, user id, users’ activity information from their profile (total posts since account creation), the total likes on posts, whether post likes were set to be visible to any 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. In Table 2 we report some of the descriptive statistics of the anorexia dataset we collected.

Separating Proana/Pro-recovery posts. The 72 tags we used to crawl Tumblr data indicated the presence of several anorexia recovery related tags (see Table 1). As also indicated in prior literature [28, 29], we thus observe that anorexia related content on Tumblr surfaces from two distinctive communities — the pro-anorexia community and the pro-recovery community. For our ensuing analysis, we developed a simple strategy to separate postings from these two communities, that looks at linguistic patterns in the tags. We manually inspected the tags we used to crawl, as well as the tags that co-occur with them in posts to create a rough sample of tags with a match for the regular expression “*recover*”. This gave us an initial sample of

⁴<http://www.tumblr.com/press>

Dataset Information (Posts)	
Number of posts is	55,334
Number of photo posts	30,865 (± 485.6)
Number of text posts	24,469
Number of notes on posts	1,521,990
Mean notes per post	27.51 (± 110.86)
Number of tags on posts	460,357
Mean tags per post	8.31 (± 132.94)
Dataset Information (Users)	
Number of unique users is	18,923
Mean posts per user is	2.92 (± 9.9)
Mean activity per user	7,203 ($\pm 13,706$)
Mean likes per user	2,359 ($\pm 6,303$)
Community-specific Information	
Pro-anorexia posts	44,033
Pro-recovery posts	11,301
Unique users in proana posts	16,105
Unique users in pro-recovery posts	3,202

Table 2: Basic descriptive statistics of Tumblr data.

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 3: Sample tags (out of 53) used for separating pro-recovery posts.

tags which are likely to be used by the pro-recovery community. We further inspected this list to discard ambiguous tags (e.g., “self harm recovery”, “anxietyrecovery”), as these tags do not distintively characterize anorexia. The final set of 53 pro-recovery related tags (ref. Table 3) uniquely distinguished pro-recovery content. We constructed the proana community to be the posts without these tags. This filtering step was applied to all of the 55,334 posts. Table 2 reports statistics on the two communities.

Control set. For the purposes of statistical validation and the prediction task in RQ 2, we collected another set of Tumblr posts to construct a control set. These are posts which do not relate to anorexia or associated disorders. We focused on the 10 most frequent Tumblr tags (“GIF”, “LOL”, “fashion”, “art”, “vintage”, “illustration”, “landscape”, “animals”, “food”, “film”)⁵. Focusing on highly used tags enables us construct a control dataset which likely captures ‘average’ Tumblr use. Over the same time period as the anorexia content crawl, we used the API to crawl ten posts for each of these frequent tags, thus giving us a set of 100 control posts per day. Our final control set consisted of 32,000 posts.

4. MEASURES

To characterize and distinguish the proana and pro-recovery communities on Tumblr, we leverage a variety of language based affective, social, cognitive, and linguistic style measures extensively examined in prior research. Our motivation springs from the fact that the psychological study of language use has received increasing scientific attention [4]—showing that it can reliably capture diagnostic information about a wide range of psychological phenomena.

Affective Processes: We propose three measures of affect: positivity, negativity, and objectivity as explored ex-

⁵<http://brandongaille.com/top-10-most-popular-tumblr-tags-and-more-tumblr/-statistics/>

tensively in text and opinion mining literature [20]. Here, an objective term is a term without any positive or negative polarity. For the purpose, we make use of the widely used and validated lexical resource SentiWordNet (<http://sentiwordnet.isti.cnr.it/>). SentiWordNet works at the synset (cognitive synonym) level as given in the WordNet database⁶, and assigns positivity, negativity, and objectivity scores to each synset.

Social Processes: We consider a variety of social processes of the proana and pro-recovery communities. Prior literature indicates that these measures reflect the extent of social engagement, nature of interactions, and social support of individuals [4]. We consider the following community-centric measures: (1) proportion of photo posts, since prior literature indicates proana to have a strong visual expression component [29]; (2) proportion of link bearing posts, that are not photos. Leveraging the various social features provided by the Tumblr website, we further propose the following measures: (3) proportion of users who make ‘likes’ on their profile posts publicly visible; (4) proportion of users whose profile allows NSFW (“not suitable for work”) content share; (5) proportion of users who allow question asking by another Tumblr user; and (6) proportion of users who allow questions from anonymous users.

We present four other measures of social processes: (1) mean number of tags per post; (2) mean notes per post; (3) mean user activity i.e., total posts since Tumblr account creation; and (4) mean number of profile ‘likes’ per user.

Cognitive Processes: For the measures belonging to cognitive processes, we largely focus on selected variables from the psycholinguistic lexicon LIWC (<http://www.liwc.net/>). We consider six types of measures of linguistic processes: (1) Interpersonal Awareness and Focus: comprising words that are *1st person singular*, *1st person plural*, *2nd person*, and *3rd person* pronouns. (2) Cognition and Perception: comprising *cognitive mech*, *inhibition*, *see*, *hear*, *feel*, *percept*, and *relative*. (3) Strengthening Cues: consisting of *negation* and *certainty*. (4) Regulatory Cues: spanning words around *non-fluency*, *insight*, *causation*, *discrepancies*, *tentativeness*. (5) Disambiguating Cues: words around *quantitative*, *inclusion*, and *exclusion*. (6) Social/Personal Concerns: words belonging to *family*, *friends*, *social*, *work*, *health*, *humans*, *religion*, *bio*, *body*, *money*, *achievement*, *home*, and *sexual*.

Linguistic Style Processes: Finally, we consider two types of measures of linguistic style [23]: (1) Lexical Density: words that are *verbs*, *auxiliary verbs*, *nouns*, *adjectives*, and *adverbs*. (2) Verbal Fluency: we consider two measures, the length of a post, and the number of words per sentence.

5. RQ 1: CHARACTERISTICS

5.1 Contrasting Tag Usage

We observe that there is relatively higher proportion of proana content (44,033 posts; 16,105 unique users) versus pro-recovery content (11,301 posts; 3,202 unique users). This demonstrates that use of Tumblr as a platform for proana support is far more widespread than for pro-recovery support. The nature of tags used by the two communities (ref. Table 4) contrasts their goals and intent of Tumblr use.

⁶WordNet is a popular and large lexical database of English language. It groups nouns, verbs, adjectives and adverbs into sets of cognitive synonyms (synsets), each expressing a distinct concept.

However, tag usage patterns of the communities shows an interesting tendency wherein the pro-recovery community, though smaller in size, makes efforts to reach out to the proana community and attempts to change their attitudes. We note that for the pro-recovery posts (total 90,298 tags), as many as 59.4% co-appear in posts on proana. This shows that the pro-recovery community tends to use tags in a way that their posts would show up in searches of proana tags on Tumblr. Through such “infiltration”, the community shows its motivation to provide help and resources to change perceptions about stereotypical idealized body image in the proana community.

On the other hand, we find that 80.98% of tags that appear in proana posts (total of 370,059 tags), uniquely *only* appear in proana posts. This demonstrates a tendency of pro-anorexics to choose and define an identity through tag use, that reinforces their perceptions of anorexia as a lifestyle.

Proana		
hip bones	thygap	flat abs
nofood	starve	must be thin
perfectbody	donteat	drugs
uglyfuck	broken inside	still fat
deadinside	selfinjury	likebarbie
Pro-recovery		
foodporn	I’m trying	selflove
body peace	never give up	togetherwecann
gain weight	good eating	morningsnack
nourishnotpunish	you are beautiful	group therapy
eatforhealth	winning	healing

Table 4: Tag usage by proana and pro-recovery communities, sorted in terms of highest Term-Frequency-Inverse Document Frequency (tf-idf).

5.2 Social Processes

Given the observed differences in the nature of tags shared by proana and pro-recovery communities, we now present distinguishing attributes manifested in the behavior of the two communities. We observe that the proana community tends to share considerably more image or link bearing content (61% vs. 37% for photos, 14% vs. 10% for links). An inspection of the tags associated with photo posts reveals a heightened focus on the self—“me”, “selfie”, “mybody”, “cheek bones”. The proana community also shows a greater propensity to allow sharing of NSFW content ($z = 7.3; p < .01$): In fact from the community’s tag usage we observe that some of the content bears the label “trigger warning” (i.e., content around self-harm, self-injury, or graphic depiction of these actions, thoughts of suicide), which is *three times higher* than that of the pro-recovery posts.

Additionally the proana community manifests an overly extrovert attitude compared to the pro-recovery community [15]. This shows in their desire to allow ‘likes’ on their posts to be visible to the public ($z = 9.6; p < .01$), as well as allowing anonymous users direct their questions to them on their Tumblr profile ($z = 13.1; p < .01$).

Next, proana content receives lower audience engagement in terms of the number of notes received per post ($z = -32.8; p < .001$), as well as in the form of ‘likes’ ($z = -14.2; p < .01$). Proana users also tend to be less active in terms of sharing content on Tumblr—an aspect that aligns with prior literature indicating the link between lowered social media use and behavioral health concerns [6].

	PA	PR
Photos	60.5%	37.1%
Links	14.2%	9.9%
Likes on posts publicly visible	80%	76.46%
Profile allows NSFW content share	0.31%	0.24%
Allow question asking	82.57%	90.74%
Allow anonymous questions	72.42%	82.75%

Table 5: Community-centric social attributes comparing proana (PA) and pro-recovery posts (PR).

	PA	PR	<i>z</i>	<i>p</i>
Tags per post	8.404	9.991	-13.1	*
Notes per post	9.212	32.166	-32.8	**
Mean user activity	6508.28	8017.32	-20.4	**
Mean profile likes	2229.36	2604.53	-14.2	*

Holm-Bonferroni correction:
* - $p \leq .01$; ** - $p \leq .001$; *** - $p \leq .0001$

Table 6: Comparison of proana (PA) and pro-recovery (PR) communities along attributes of social processes⁸.

5.3 Affective and Linguistic Style Processes

In the light of measures of affective processes (Table 7) pro-anorexics are observed to be less positive ($z = -19.8$; $p < .001$), more negative ($z = 27$; $p < .001$), and less objective in contrast to recovering anorexics (based on Mann-Whitney U tests). In essence, compared to the proanorexics, the pro-recovery community exhibits a more pronounced hedonic focus on positive emotions and an objective outlook towards life. The negativistic tone of proana content reflects a lowered sense of self-esteem, likely due to dissatisfaction with normative notions of idealized body weight and shape [27].

Proana content contains more verbs ($z = 25.3$; $p < .001$), auxiliary verbs, and adverbs ($z = 59.1$; $p < .0001$) (which indicate discourse around actions), based on Mann-Whitney U tests. But it contains higher mentions of nouns and adjectives ($z = 35.1$; $p < .001$), which reveal heightened interest in objects and things around them [4]. Greater action-focused expressions might be indicative of their repeated reinforcement of proana-related lifestyle choices and activities, intended or taken to maintain this routine consistently. Proana content also bears greater mentions of swear words ($z = 19.6$; $p < .001$)—a known sign of aggression [5], as well as more attributes of non-fluency ($z = 47.1$; $p < .0001$) including lower verbal fluency (post length and words/sentence are both lower compared to pro-recovery posts). This might point to their negative attitude, informal nature of discourse and lowered cognitive complexity [1].

5.4 Cognitive Processes

From Table 9, we find that proana content contains significantly fewer cognition and perception words (category mean value less than half compared to the pro-recovery community), based on Mann-Whitney U tests. This is known to be a correlate of cognitive impairment and emotional instability [5]. In fact, it is known that cognitive function is impaired in individuals who are at extremely low

⁸Statistical comparisons ($df = 55332$) are based on Mann-Whitney U test and familywise error rate is controlled for using Holm-Bonferroni correction (* - $p \leq .01$; ** - $p \leq .001$; *** - $p \leq .0001$)

	PA	PR	<i>z</i>	<i>p</i>
PA	0.2798	0.2939	-19.8118	**
NA	0.1972	0.1873	27.0454	**
objectivity	0.4062	0.4733	-10.2583	-

Holm-Bonferroni correction:
* - $p \leq .01$; ** - $p \leq .001$; *** - $p \leq .0001$

Table 7: Comparison of proana (PA) and pro-recovery (PR) communities along attributes of affective processes⁷.

	PA	PR	<i>z</i>	<i>p</i>
Lexical density				
verbs	0.0037	0.002	25.33	**
aux_verbs	0.0739	0.0599	53.32	***
nouns	0.0247	0.024	10.11	-
adjectives	0.0070	0.0049	35.11	**
adverbs	0.0583	0.0235	59.05	***
swear	0.0245	0.0116	19.637	**
nonfluency	0.1025	0.0850	47.15	***
<i>Category mean</i>	<i>0.0358</i>	<i>0.0378</i>		
Verbal fluency				
post length	59.662	99.459	-40.59	***
words/sentence	11.294	17.352	-19.23	*

Holm-Bonferroni correction:
* - $p \leq .01$; ** - $p \leq .001$; *** - $p \leq .0001$

Table 8: Comparison of proana (PA) and pro-recovery (PR) communities along attributes of linguistic style processes⁷.

body weight [15]. Proana users also exhibit higher self-attentional focus as indicated by greater use of first person pronouns ($z = 33.1$; $p < .0001$). They are less socially interactive as is shown in lowered use of second person pronouns ($z = -42.9$; $p < .0001$); but pay greater attention to people and objects. Higher self-occupation is a known correlate of behavioral health concerns [6].

Proana community further uses fewer strengthening ($z = -42.1$; $p < .0001$ for “negation”), regulatory ($z = -50.4$; $p < .0001$ for “causation”), and disambiguating cues ($z = -39.6$; $p < .0001$ for “exclusion”) in their discourse on Tumblr. This shows an avoidant attitude, and a lack of honesty in their conversational tone [4]. The findings also demonstrate that the proana individuals manifest fewer degree of differentiation, articulation, and integration within their cognitive systems. Finally, proana community uses fewer words related to social and personal concerns ($z = -44.4$; $p < .0001$ for the words in “social”) compared to the recovering group (similar findings in [15]). Thus the former community shows little evidence about being socially or professionally embedded with their family or friends (note the word sets “family”, “friends”, “home”, “work”, “social”). This could be due to feelings of social isolation and rejection, or the lack of social support.

To summarize, the proana community indicates a more concentrated focus on the present circumstances, reduced level of cognitive processing, avoidant attitude and behavior, and a heightened degree of self-preoccupation. On the other hand, likely due to their realization of the ill effects of anorexia, pro-recovery community depicts a greater tendency to be ruminative and past-oriented. The recovery course also lets them gather greater insight about their illness. Hence they exhibit greater cognitive processing. Further, they show lowered focus on self, and the heightened

	PA	PR	z	p
Interpersonal awareness				
1stpp	0.0358	0.0293	33.14	***
2ndpp	0.0216	0.0407	-42.97	***
Cognition and Perception				
cognitive	0.0173	0.0237	-22.88	**
inhibition	0.0349	0.0515	-49.65	***
see-hear-feel	0.0072	0.0100	-58.96	***
perception	0.0037	0.0050	-34.73	**
relative	0.0156	0.0271	-53.73	***
<i>Category mean</i>	<i>0.0157</i>	<i>0.0235</i>		
Strengthening cues				
certainty	0.0194	0.0277	-13.49	*
negation	0.0112	0.0173	-42.14	***
<i>Category mean</i>	<i>0.0153</i>	<i>0.0225</i>		
Regulatory cues				
causation	0.0334	0.05015	-50.44	***
insight	0.0377	0.0452	-23.6	**
<i>Category mean</i>	<i>0.0628</i>	<i>0.0736</i>		
Disambiguating cues				
exclusion	0.0174	0.0246	-39.63	**
quantitative	0.0127	0.0183	-21.61	**
<i>Category mean</i>	<i>0.1193</i>	<i>0.1365</i>		
Social/personal concerns				
achievement	0.0048	0.0074	-47.46	***
family	0.0134	0.0266	-17	*
friends	0.0394	0.0552	-33.89	**
home	0.0002	0.0126	-54.83	***
social	0.0033	0.0094	-44.44	***
sexual	0.3531	0.2506	41.5	***
work	0.1185	0.2807	-32.35	**
<i>Category mean</i>	<i>0.0829</i>	<i>0.0917</i>		

Holm-Bonferroni correction:

* - $p \leq .01$; ** - $p \leq .001$; *** - $p \leq .0001$

Table 9: Comparison of proana (PA) and pro-recovery (PR) communities along attributes of cognitive processes⁷. (Only significant measures listed.)

use of future tense words tells us that they are goal-oriented and look forward to the recovery process in the near future.

6. RQ 2: PREDICTION

Finally, we pursue the use of supervised learning to construct classifiers trained to predict: (1) whether a post shares any kind of anorexia related content; and (2) whether a post relates to the proana or the pro-recovery community.

To this end, we use observations of our four categories of measures as individual features of the prediction model: social, affective, linguistic style, and cognitive processes. Each Tumblr post is represented as a vector of these features. For evaluation purposes, we use standard 10-fold cross validation (CV) and 1000 random 10-fold CV assignments for generalizability of the results. We train a binary Support Vector Machine (SVM) classifier corresponding to each task (best performing classifier among a set of seven classifiers that were tested). We report the average accuracy, precision, recall, F1, specificity, and AUC (area under curve) as metrics of performance. To understand the performance of each category of measures, we train and test five different models per prediction task—four models corresponding to the four measure categories: social, affective, cognitive, and

	acc	prec	rec	F1	spec	AUC
affective	59.9	0.582	0.528	0.554	0.552	0.539
social	63.3	0.623	0.644	0.658	0.668	0.638
linguistic	71.2	0.708	0.742	0.712	0.699	0.712
cognitive	74.2	0.719	0.701	0.703	0.732	0.722
all	81.4	0.844	0.834	0.818	0.797	0.794

Table 10: Predicting proana versus pro-recovery.

linguistic style processes; and the fifth model using all measure categories.

Anorexia Vs. Control. For the first task, we intend to distinguish between the following two classes: the first set comprising the combined set of proana and pro-recovery posts, and the second comprising the control set posts (ref. section on “Data”). We use balanced classes for appropriate model training, which yields us 32,000 posts for each class. Prediction results indicate that we are able to detect if a post shared anorexia related content (versus not) with 83% accuracy (33% more over baseline performance of 50%) and .85 precision, when using all measure categories as features. The model using measures of cognitive processes alone bears the most predictive power among the set of all measures categories (accuracy = 74%; F1 = .76).

Proana Vs. Pro-recovery. For the second prediction task, our binary SVM classifier differentiates between the proana and the pro-recovery sets of posts. We again use balanced class sizes i.e., 11,301 posts per class. The prediction results for this task are shown in Table 10. The best performing model is the fifth model, i.e., the one that uses all measures (accuracy = 81.4%; F1 = .81). Among the four measure categories, the performances align with our findings from RQ 1: best performance is given by the cognitive processes model (accuracy = 74%; F1 = .7), whereas the affective processes model yields the poorest performance (accuracy = 59%; F1 = .55). All models, however show improvement over baseline (baseline accuracy = 50%).

Summarily, our results indicate that Tumblr use related to anorexia bears distinctive markers in terms of social, affective, cognitive, and linguistic style expression. This can be used to detect anorexia related posts in an automated setting. Further, the two prominent communities: proana and pro-recovery, themselves exhibit markedly differentiating characteristics which can be used to further identify the kinds of anorexia related content shared on Tumblr.

7. DISCUSSION

7.1 Clinical Relevance

Our findings reveal several characteristics of social media (Tumblr) use for anorexia. Many of our findings on Tumblr use for pro-anorexia and pro-recovery align with prior literature [15, 27, 28, 9, 3]. Together, they demonstrate that although Tumblr is a new social media, the broad purposes of the use of this platform for anorexia and recovery are similar to other online platforms. However, Tumblr being a multimedia-rich microblogging platform, we observe that proana content that is shared is often replete with triggering and graphic images, aside from text. Further, our data indicates that pro-anorexia content on Tumblr is notably more pervasive compared to pro-recovery content (almost

five times in terms of unique users). Weaving together these observations, it appears that, in contrast with other online technologies examined in prior literature [3, 9, 15, 27, 28], Tumblr is used disproportionately more frequently as a support platform on maintenance of anorexic lifestyle.

Consequently, concrete steps must be taken to change individuals' attitudes or beliefs towards self-hatred or misperceptions about idealized physical appearance. The pro-recovery community attempts to correct these attitudes by permeating into the proana community, and Tumblr also issues public service announcements on first-time pro-anorexia related tag searches from a user account. However the large presence of proana community on Tumblr shows that more effective methods of intervention are needed to bring help and support to affected populations in a timely manner. We have shown that analysis of large-scale social media activity may reveal distinctive affective, social, cognitive, and linguistic style characteristics that could surface such content. From a clinical perspective, our methods could complement self-report based psychiatric measurements of eating disorders. They can also help tailor interventions through non-reactive and non-intrusive assessment of behavior of proana populations, as discussed below.

Proana interventions. Tumblr already has a content policy in place, which directs a user attempting to search on pro-ED or other self-harm content on the website to a public service announcement page with the following message:

Everything okay? If you or someone you know is struggling with an eating disorder, NEDA is here to help: call 1-800-931-2237 or chat with them online. If you are experiencing any other type of crisis, consider talking confidentially with a volunteer trained in crisis intervention at www.imalive.org, or anonymously with a trained active listener from 7 Cups of Tea. [...]

We envision new kinds of provisions to be added on top of current efforts like this. The goal of such provisions would be to actively discourage or prevent individuals from seeking or sharing proana content: (1) Searches on proana content may automatically direct individuals to links hosting helpful and research-supported resources, highlighting the health risks of anorexia (e.g., the National Eating Disorders Association NEDA website¹). (2) Our methods could be used to detect whether a Tumblr user is attempting to post anorexia promoting or other triggering content. At that point the system could intercept with a pop-up message that provides link to an eating disorder helpline. Additionally, on detecting proana posts with unusually high negative affect, lowered cognitive processing, increased isolation or depressive tendencies, social media sites could send a private message to the user about connecting anonymously to a trained crisis intervention volunteer.

Pro-recovery interventions. Design considerations for the pro-recovery community include provisions that make Tumblr (or another social media platform) conducive to their recovery journey: (1) Prior literature has demonstrated that the web can be a wonderful resource for those seeking help and support with mental health concerns [21]. Thus, recommendations could be surfaced on one's Tumblr dashboard: e.g., the system could match up recovering anorexics so that they can receive and provide peer-motivation.

(2) Pro-recovery Tumblr users may be provided with diary-centric tools on the website which would allow them to log artifacts (e.g., posts, images, links) they share on Tumblr along their recovery journey. Such aggregated logs, if desired and with consent, may be shared with a caregiver or a clinician time to time, so that they can track the trajectory of their recovery process and make therapeutic accommodations, if needed.

7.2 Ethical Considerations and Limitations

While thinking about design provisions, it is important to bear in mind that anorexia content can be extremely sensitive and often comes from vulnerable and emotionally fragile communities who feel uncomfortable disclosing such information in offline or mainstream settings. We would like to emphasize that design considerations in this space need to honor the privacy of the affected individuals and abide by appropriate ethical guidelines that ensure that the intended benefits of these interventions outweigh the risks. This can, for instance, be done by revealing social media detected risk to either the individual themselves, or to an identified trusted social contact or clinician. Summarily, we hope this line of research not only opens up new avenues to detect and help individuals sharing anorexia content, but also sparks interaction with ethics and eating disorder support and clinician communities.

It is worthwhile noting that our paper does not make any claim as to attributing anorexia to the posters of the Tumblr content we study. We acknowledge that it is not clear to what extent the individuals we studied actually met clinical DSM criteria on anorexia. We also note that while our predictive framework was able to detect anorexia related content with high accuracy, nevertheless it is prone to false positives/negatives. We caution against using these methods as standalone techniques for diagnosis. We believe our methods would be most successful with involvement of clinicians or psychiatrists who work with pro-anorexia communities.

Finally, our findings also raise difficult questions for health researchers about resistance to treatment, anti-recovery positions, and the general self-loathing notion associated with anorexia. Are such social media content a lethal threat to vulnerable young people? How effective would the proposed interventions be, in terms of suppressing pro-anorexia content? Or will such content resurface in a different guise, a common characteristic of most deviant communities?

8. CONCLUSION

Online pro-anorexia content is a case of societal concern and many health authorities and organizations have been attempting to curb its effects. In this paper, we leveraged the observation that most individuals affected by this disorder fall within the age group in which social media platforms such as Tumblr are used heavily. Our results show that, compared with recovering anorexics, pro-anorexics expressed greater negative affect and lesser positive affect, lowered cognitive complexity in their language, increased aggregation and self-occupation, and greater feelings of social isolation. We leveraged these distinctive characteristics in a predictive setting to identify proana or pro-recovery content with high accuracy. Our findings bear implications in the design of social media interventions that can bring in-time help and support to anorexia communities.

9. REFERENCES

- [1] Jaime Arguello, Brian S Butler, Elisabeth Joyce, Robert Kraut, Kimberly S Ling, Carolyn Rosé, and Xiaoping Wang. Talk to me: foundations for successful individual-group interactions in online communities. In *Proc. CHI*, pages 959–968. ACM, 2006.
- [2] Anna M Bardone-Cone and Kamila M Cass. Investigating the impact of pro-anorexia websites: A pilot study. *European Eating Disorders Review*, 14(4):256–262, 2006.
- [3] Dina LG Borzekowski, Summer Schenk, Jenny L Wilson, and Rebecka Peebles. e-ana and e-mia: A content analysis of pro-eating disorder web sites. *American journal of public health*, 100(8):1526, 2010.
- [4] Cindy Chung and James W Pennebaker. The psychological functions of function words. *Social communication*, pages 343–359, 2007.
- [5] J David Creswell, Suman Lam, Annette L Stanton, Shelley E Taylor, Julianne E Bower, and David K Sherman. Does self-affirmation, cognitive processing, or discovery of meaning explain cancer-related health benefits of expressive writing? *Personality and Social Psychology Bulletin*, 33(2):238–250, 2007.
- [6] Munmun De Choudhury, Scott Counts, and Eric Horvitz. Predicting postpartum changes in emotion and behavior via social media. In *Proc. CHI*, pages 3267–3276. ACM, 2013.
- [7] Munmun De Choudhury, Scott Counts, Eric Horvitz, and Aaron Hoff. Characterizing and predicting postpartum depression from facebook data. In *Proc. CSCW*. ACM, 2014.
- [8] Karthik Dinakar, Birago Jones, Henry Lieberman, Rosalind Picard, Carolyn Rose, and Matthew Thoman. You too?! mixed initiative lda story-matching to help teens in distress. In *Proc. ICWSM*, 2012.
- [9] Jeff Gavin, Karen Rodham, and Helen Poyer. The presentation of “pro-anorexia” in online group interactions. *Qualitative Health Research*, 18(3):325–333, 2008.
- [10] Hans Wijbrand Hoek. Incidence, prevalence and mortality of anorexia nervosa and other eating disorders. *Current opinion in psychiatry*, 19(4):389–394, 2006.
- [11] Christopher M Homan, Naiji Lu, Xin Tu, Megan C Lytle, and Vincent Silenzio. Social structure and depression in trevorspace. In *Proc. CSCW*, pages 615–625. ACM, 2014.
- [12] Sue Jamison-Powell, Conor Linehan, Laura Daley, Andrew Garbett, and Shaun Lawson. I can’t get no sleep: discussing# insomnia on twitter. In *Proc. CHI*, pages 1501–1510. ACM, 2012.
- [13] Adrienne S Juarascio, Amber Shoaib, and C Alix Timko. Pro-eating disorder communities on social networking sites: a content analysis. *Eating disorders*, 18(5):393–407, 2010.
- [14] Amanda Lenhart, Kristen Purcell, Aaron Smith, and Kathryn Zickuhr. Social media and young adults. *Pew Internet & American Life Project*, 3, 2010.
- [15] Elizabeth J Lyons, Matthias R Mehl, and James W Pennebaker. Pro-anorexics and recovering anorexics differ in their linguistic internet self-presentation. *Journal of Psychosomatic Research*, 60(3):253–256, 2006.
- [16] Ruaidhri Mulveen and Julie Hepworth. An interpretative phenomenological analysis of participation in a pro-anorexia internet site and its relationship with disordered eating. *Journal of health psychology*, 11(2):283–296, 2006.
- [17] Elizabeth L Murnane and Scott Counts. Unraveling abstinence and relapse: smoking cessation reflected in social media. In *Proc. CHI*, 1345–1354. ACM, 2014.
- [18] Mark L Norris, Katherine M Boydell, Leora Pinhas, and Debra K Katzman. Ana and the internet: A review of pro-anorexia websites. *Intl Journal of Eating Disorders*, 39(6):443–447, 2006.
- [19] Hyun Jung Oh, Carolyn Lauckner, Jan Boehmer, Ryan Fewins-Bliss, and Kang Li. Facebooking for health: An examination into the solicitation and effects of health-related social support on social networking sites. *Computers in Human Behavior*, 29(5):2072–2080, 2013.
- [20] Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. Thumbs up?: sentiment classification using machine learning techniques. In *Proc. ACL*, pages 79–86. Association for Computational Linguistics, 2002.
- [21] Minsu Park, David W McDonald, and Meeyoung Cha. Perception differences between the depressed and non-depressed users in twitter. In *Proc. ICWSM*, 2013.
- [22] Michael J Paul and Mark Dredze. You are what you tweet: Analyzing twitter for public health. In *Proc. ICWSM*, 2011.
- [23] James W Pennebaker, Tracy J Mayne, and Martha E Francis. Linguistic predictors of adaptive bereavement. *J. personality and social psychology*, 72(4):863, 1997.
- [24] Sarah Swan and Bernice Andrews. The relationship between shame, eating disorders and disclosure in treatment. *British journal of clinical psychology*, 42(4):367–378, 2003.
- [25] 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, et al. Misleading health-related information promoted through video-based social media: anorexia on youtube. *J. medical Internet research*, 15(2), 2013.
- [26] Kelly Bemis Vitousek, Jennifer Daly, and Christopher Heiser. Reconstructing the internal world of the eating-disordered individual: Overcoming denial and distortion in self-report. *Intl Journal of Eating Disorders*, 10(6):647–666, 1991.
- [27] Markus Wolf, Jan Sedway, Cynthia M Bulik, and Hans Kordy. Linguistic analyses of natural written language: Unobtrusive assessment of cognitive style in eating disorders. *Intl journal of eating disorders*, 40(8):711–717, 2007.
- [28] Markus Wolf, Florian Theis, and Hans Kordy. Language use in eating disorder blogs psychological implications of social online activity. *J. Language and Social Psychology*, 32(2):212–226, 2013.
- [29] Elad Yom-Tov, Luis Fernandez-Luque, Ingmar Weber, and Steven P Crain. Pro-anorexia and pro-recovery photo sharing: A tale of two warring tribes. *J. medical Internet research*, 14(6), 2012.