Fat-Burning or Fat-Shaming? Theme and Community Overlap Between Online, Weight-Stigmatizing and Exercise-Promoting Networks.

Ahuitz Rojas-Sánchez (✉ Ahuitzrojas@outlook.com)

Universidad Autonoma de San Luis Potosí <https://orcid.org/0000-0001-6672-2100>

Jenine K. Harris

Washington University in St. Louis

Philippe Sarrazin

Université Grenoble Alpes

Aïna Chalabaev

Université Grenoble Alpes

Research Article

Keywords: Twitter, Weight stigma, Social Networks, Exercise, Social Media.

Posted Date: May 24th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-526129/v1>

License: This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Abstract

Purpose: This study aimed to determine if networks of users consistently posting about exercise and fat exist and overlap on social media sites.

Method: We collected 3,772,507 posts from Twitter that included the words “fat” and “exercise”. Using network structure methods, we identified communities of interconnected users and overlaps between those tweeting “fat” and those tweeting “exercise”.

Results: Common word pairings were identified using Natural Language Processing (NLP). Networks of users consistently talking about exercise ($n=3,573$) and fat ($n=2,007$) were found on Twitter. An increased mean total-degree and reduced average path length indicate that the fitness-talk network serves as a connecting bridge between highly scattered communities of the weight-talk network.

Conclusion: We identified groups on Twitter dedicated to consistently producing weight stigmatizing content and promoting exercise with weight-loss messages. These groups partially overlap with pro-health groups which could lead to users looking for exercise advice in Twitter to find themselves immersed in a stigmatizing network.

Introduction

Around 90% of posts on social networking sites (SNS) containing the words “fat” or “obesity” are found on Twitter, while Facebook, blogs, forums and internet comments contain the other 10% [1]. Out of posts containing the words “fat”, “obesity”, and “overweight”, the word “fat” is by far the most common [1]. This makes Twitter the most important place to study online social networks, and the word “fat” the keyword to use when looking for information related to weight. Furthermore, posts containing the word “fat” have been found to have a mostly negative connotation. For example, one study found that 56% of the messages containing the word “fat” were explicitly negative towards higher weight individuals [2]. This suggests the presence of weight stigma, the pervasive social devaluation and denigration of overweight and obese people, which may be problematic given that weight stigma is related to unhealthy behaviors, such as high calorie intake, binge eating, and physical inactivity [3]. Specifically, negative portrayals of obesity in the media have been found to increase caloric intake and reduce physical activity and dieting self-efficacy [4].

Communities promoting weight loss online do so while overlapping with those promoting fitness [5, 6], which means that users looking for information related to healthy exercise habits might inadvertently end up receiving stigmatizing content. This could happen because weight loss messages are common within twitter stigmatizing communications [2]. Alternatively, the word “fat” works both as an adjective with a negative connotation (i.e., fat person), and as a noun that is commonly associated with exercise (i.e., burn fat), which could result in a bundling together of both concepts by Twitter’s recommendation algorithm. Finally, there is large body of research suggesting that individuals have an inaccurate understanding of

physical activity, mainly reducing it to weight loss and appearance change [7], which could result in an overlap between stigmatizing and fitness promoting content.

The focus on weight when promoting exercise could be particularly prejudicial if higher weight users are the ones searching for exercise-related information. Indeed, these users could be exposed to stigmatizing content if there is an overlap in themes and users. This study aimed to determine if networks of users consistently talking about exercise and fat exist on Twitter (objective 1). Assuming they do, we hypothesized that they overlap (objective 2), and that, in accordance with previous findings, the content of their communications will be negative, or weight loss-related (objective 3).

Methods

Data collection

Data was collected using the Twitter Application Programming Interface (API), a web-based program that allows users to interact with Twitter's data (i.e., tweets and metadata). This allowed for keyword searches of all tweets containing the words "exercise" and/or "fat". These searches were implemented daily over a period of 3 months (November 2017 to February 2018). Each daily search provided the most recent tweets of the day (up-to 50,000). It also provided information on how many times each tweet was retweeted, the user-id of the user who produced the tweet, and the number of followers of that user.

Community structure and overlap

In order to determine if networks of users consistently talking about exercise and fat exist on Twitter (objective 1) users who tweeted at least once a week about "fat" or "exercise" were assigned to a core group, while those who only tweeted once in three months about "fat" or "exercise" were assigned to the visiting group. Users who tweeted more than once but not once a week were excluded from the study.

In order to study a network of Twitter users, each user is considered as a node. A node can follow another node, creating a link (e.g., $A \rightarrow B$). In $A \rightarrow B$, node A has an out-degree of 1 (because it follows B) and node B has an in-degree of 1 (because it is followed by A). In social networks, the direction of the arrow typically represents information flow, but in the case of Twitter the relationship is inversed because users follow other users to see the content they publicly post, not to send direct information. In short, if node A follows node B ($A \rightarrow B$), node A is seeing information posted by node B.

Nodes with a large following within the network have a high in-degree, while nodes following a large number of users within the network have a high out-degree, with the sum of both hereby referred to as total-degree. Additionally, within the social network, communities can emerge [8]. These communities are composed of users with more connections among them but fewer connections to nodes outside the community which may themselves be part of different communities.

Relationships between users were mapped to allow for social network analysis and to determine network structure. The user-id from the core group was used to identify the followers of each individual user. The relationships were then filtered to only include relationships within the core group. Using these relationships, a core weight-talk network and a fitness-talk network were mapped. Additionally, to determine if there is an overlap between weight-talk and fitness-talk communities, a combined network was mapped, and the metrics of this network compared to the previous ones.

To answer questions regarding structure and overlap (objective 2), four standard network metrics were used: density, average path length, mean total-degree, and clustering coefficient. Density measures how many connections exist in the network and ranges from zero (no connections) to one (all possible connections exist). Average path length refers to the average shortest paths between two nodes and is important because it indicates how far in average, information has to travel to reach from one node to another. Mean total-degree refers to the average number of connections per node and is useful for determining how well connected a network is. Finally, clustering coefficient is a measure of the extent to which nodes in a network tend to cluster together, which allows us to understand the network structure in reference to its communities.

Additionally, modularity was calculated within the network. This is the identification of communities within a network based on the similarities between their connections. It was done using the fast unfolding of communities in large networks algorithm [8]. This algorithm decomposes the networks into sub-units or communities, which are sets of highly inter-connected nodes. Four main communities per network were extracted. Hubs were found by ranking the nodes based on their in-degree. This allowed for a linguistic corpus analysis based on the communities of each network.

Linguistic corpus

The linguistic corpus is the whole set of text data (tweets) to be analyzed. Simple Natural Language Processing (NLP) techniques like the division of sentences into individual words (tokenization) and their subsequent analyses in duos or triplets (n-grams) were used to analyze this data. Latent Dirichlet Allocation (LDA), which allows for the grouping of observations (words) to be explained by unobserved groups (themes), was also used. A total of 3,772,507 tweets were collected, however out of those, non-English-language tweets (n=510,145) and tweets from people that were not in the core or visiting categories (n=1,291,155) were removed. As a result, the total corpus was reduced to 1,971,207 tweets.

Linguistic n-grams

In order to confirm that the communicational content of the weight-talk and fitness-talk network is explicitly negative, or weight loss related (objective 3) the tweets were divided into the four main communities of each core social network. After that, a list of linguistic bigrams and trigrams that excluded prepositions, conjunctions, and linguistic fillers was generated. Finally, a Latent Dirichlet

Allocation (LDA) model was applied to differentiate individual unobserved clusters of words, which allowed for a qualitative lexical text analysis and extraction of common themes within the communities.

Results

Network characteristics

The weight-talk and fitness-talk networks were similar in size, both in its core ($N_{\text{fat}}=2,007$; $N_{\text{exercise}}=3,573$) and visiting forms ($N_{\text{fat}}=737,102$; $N_{\text{exercise}}=703,948$). The number of tweets, however, was smaller in the core weight-talk network. A total of 3,573 users consistently use their profiles as a way to communicate about exercise, while 2,007 use them for “fat” related communication, indicating the existence of dedicated networks. The number of users and tweets in the visiting and core networks, as well as the four main communities, are described in Table 1.

Network structure

The structure and communities of the weight-talk and fitness-talk networks was calculated using the fast unfolding of communities in large networks algorithm [8]. The mean total-degree is larger in the fitness-talk network (13.40) than in the weight-talk one (5.54) indicating that the fitness-talk network users are more connected overall, which makes the average path length smaller (3.79 for the fitness-talk, and 6.02 for the weight-talk network). The communities within the networks, however, seem to be more clustered together in the weight-talk network (0.17) than in the fitness-talk one (0.093), creating a divided network with a larger diameter (24 for the weight-talk network compared to 16 in the fitness-talk one).

Overlap between networks

A small percentage (7.6%) of users are in both networks. The mean total-degree of the combined network (12.91) is larger than the averaging of the means of the individual networks (9.71). The average path length (4.31), and the clustering coefficient (0.107) for the combined network drop below the averages of both individual networks, while the diameter is widely reduced from the weight-talk network. This could indicate that the users in the fitness-talk network serve as a connecting bridge between highly scattered communities of the weight-talk network. Table 2 shows descriptions of the networks.

Content of communications

To understand the overall content of each group, the tweets were divided into individual words (tokenization), pairs of words (bigrams) or triplets of words (trigrams). The lists reveal themes of physical fitness, health, and fitness for health in the fitness-talk network’s communities. As for the weight-talk

network, the communities contain mainly words regarding weight loss or fat burning, insults, and blame regarding food and eating behaviors.

Discussion

The belief that health can be achieved through weight loss reduction behaviors such as physical activity seems to be common [7]. This study aimed to test this overlap within twitter communities that mainly shared and produced content relating to weight and fitness by posing three important questions: Are there dedicated communities of users discussing “fat” and “exercise” online? If so, do these communities overlap? and what are the themes of these communities?

Results corroborated the existence of a large core network of users tweeting about fitness or weight at least once per week. Dense and highly clustered communities were found within the individual networks with the communities in the weight-talk network mainly focusing on weight loss, stigmatization, and internalized stigma, and communities in the fitness-talk network focusing on exercise for health and exercise for fitness. While the geodesic distance in the fitness-talk network was higher than that of the weight-talk network, the clustering coefficient was not. This corroborates previous findings that weight-loss promoting communities are more closely knitted together than pro-health ones [5]. It is also in line with research that has found that health-promoting tweets focus on nutrition and exercise, while pro-thinness ones focus on thinness and disordered eating behaviors [6].

The combination of both networks showed a mild overlap between them, which could indicate that the users in the fitness-talk network serve as a connecting bridge between highly scattered communities of the weight-talk network. This is consistent with previous studies showing the prevalence of pro-thin type communications among weight-talk networks in Twitter [1, 2]. Given that the Twitter algorithm recommends people whom the user is most likely to interact with, this could create an ever-growing community of bias. On the other hand, there could be an ingrained association between exercise and weight loss, through concepts like “fat-burning” [7]. This, in turn, might lead to people looking for pro-health advice to find themselves within a stigmatizing network.

Finally, this study examined what people were tweeting and sharing about fat and exercise. We found that the networks were not only divided as a function of their interconnectivity, but also by their topics. In the case of the fitness-talk network, fitness and health-related exercise tweets were found to be the most common reasons for practicing physical activity. Other topics were more generic and unrelated to physical activity, like English learning and politics. In the weight-talk network, on the other hand, the most popular topics were weight loss-related tweets. This resembled previous findings where weight loss was a common theme of Twitter [2]. The second most common category of tweets in the weight-talk network was explicit weight stigmatization, confirming previous findings [1, 2]. Internalized stigma tweets, which is to say tweets making explicit self-directed stigmatizing comments, were found in a specific sub-cluster that also happens to house pro-anorexia accounts. And finally, a cluster of porn related accounts was found that held no relation to the study.

Apart from the direct descriptive implications of the study, the results can be framed within the broader theoretical framework of Group Norms Theory [9], which states that intergroup attitudes and behaviors are strongly influenced by representatives of psychologically valued “reference” groups. The weight-talk network in particular is shown to have clustered communities within it. These communities have explicitly stigmatizing messages and pro-thinness communications that could arise contagion of these sorts of attitudes. Further study is needed to address the possible transferable nature of stigma.

This has a wide array of public health implications. Understanding the transfer and propagation of stigma could help in the design of campaigns directed at prevention (i.e., via information), inoculation (i.e., coping mechanisms), and intervention of stigmatizing attitudes in a large-scale fashion. Specifically, individual targeting to the hubs that are shared in both networks could result in a reframing of the health and exercise promotion communications [10]. Another important alternative to consider is the targeting of subcommunities, like the pro-thinness one in the weight-talk network or the pro-health one in the fitness-talk network, to try and change the rhetoric of the communications [10].

While this study has described the structure and themes of weight and fitness networks on Twitter, some important questions still remain about the influence of the content posted in Twitter on in other user's levels of physical activity, as well as the predictive capacity of the textual content produced by the users over their weight stigma attitudes or physical activity habits.

This study adds to existing knowledge in three important ways: (1) it determined the existence of a dedicated weight stigmatizing community, (2) it showed partial overlap between fitness-talk and weight-talk communities on Twitter, and (3) it confirmed previous findings regarding the content of communications within and between these communities.

Declarations

Founding: This research was funded by The CONACYT (Consejo Nacional de Ciencia y Tecnología) International Scholarship under which the corresponding author's Ph.D. was funded.

Conflicts of interest/Competing interests: The authors declare no conflict of interests.

Availability of data and material: Data is available upon request.

Code availability: Code is available upon request.

References

1. Chou WS, Prestin A, Kunath S (2014) Obesity in social media: a mixed methods analysis. *Transl Behav Med* 4:314–323. <https://doi.org/10.1007/s13142-014-0256-1>
2. Lydecker JA, Cotter EW, Palmberg AA et al (2016) Does this Tweet make me look fat? A content analysis of weight stigma on Twitter. *Eat Weight Disord - Stud Anorex Bulim Obes* 21:229–235.

<https://doi.org/10.1007/s40519-016-0272-x>

3. Tomiyama AJ, Carr D, Granberg EM et al (2018) How and why weight stigma drives the obesity 'epidemic' and harms health. BMC Med 16:123. <https://doi.org/10.1186/s12916-018-1116-5>
4. Lambert ER, Koutoukidis DA, Jackson SE (2019) Effects of weight stigma in news media on physical activity, dietary and weight loss intentions and behaviour. Obes Res Clin Pract 13:571–578. <https://doi.org/10.1016/j.orcp.2019.09.001>
5. Tiggemann M, Churches O, Mitchell L, Brown Z (2018) Tweeting weight loss: A comparison of #thinspiration and #fitspiration communities on Twitter. Body Image 25:133–138. <https://doi.org/10.1016/j.bodyim.2018.03.002>
6. Harris JK, Duncan A, Men V et al (2018) Messengers and Messages for Tweets That Used #thinspo and #fitspo Hashtags in 2016. Prev Chronic Dis 15:. <https://doi.org/10.5888/pcd15.170309>
7. Harris J, Cale L, Duncombe R, Musson H (2018) Young people's knowledge and understanding of health, fitness and physical activity: issues, divides and dilemmas. Sport Educ Soc 23:407–420. <https://doi.org/10.1080/13573322.2016.1228047>
8. Blondel VD, Guillaume J-L, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. J Stat Mech Theory Exp 2008:P10008. <https://doi.org/10.1088/1742-5468/2008/10/P10008>
9. Crandall CS, Eshleman A (2003) A justification-suppression model of the expression and experience of prejudice. Psychol Bull 129:414–446. <https://doi.org/10.1037/0033-2909.129.3.414>
10. Valente TW (2012) Network Interventions. Science 337:49–53. <https://doi.org/10.1126/science.1217330>

Tables

Table 1. *Number of users and posts by network and community.*

	Exercise	Fat
Users		
Visiting	703,948	737,102
Core	3,573 (100%)	2,007 (100%)
Core (no isolates)	3,557 (99.55%)	1,827 (91.03%)
Community #1	1,363 (38.15%)	662 (36.23%)
Community #2	1,107 (30.98%)	286 (15.65%)
Community #3	491 (13.74%)	129 (7.06%)
Community #4	475 (13.27%)	125 (6.84%)
# Tweets		
Visiting	703,948	737,102
Core	399,641	130,516
Community #1	142,448	50,613
Community #2	41,220	4,783
Community #3	11,837	2,295
Community #4	12,349	3,206

Table 2. *Descriptive statistics of the fitness-talk, weight-talk, and combined networks*

	Fitness-talk network	Weight-talk network	Combined network
Nodes	3,573	2,007	5,695
Edges (Links)	47,695	12,210	73,520
Semi-isolates	16	180	364
Nodes in giant component	3,557	1,827	5,331
Edges in giant component	47,677	11,955	71,627
Modularity (without isolates; Resolution 1.5)	0.24	0.53	0.20
Density (without isolates)	0.004	0.004	0.003
Communities in giant component	16	26	36
Average path length	3.79	6.02	4.31
Diameter	14	24	16
Mean total-degree	13.35	6.08	12.91
Mean total-degree (without isolates)	13.40	5.54	13.44
Mean clustering coefficient	0.093	0.17	0.107