Finding influential users of online health communities: a new metric based on sentiment influence

Kang Zhao,[1] John Yen,[2] Greta Greer,[3] Baojun Qiu,[4] Prasenjit Mitra,[2] Kenneth Portier[3]

[1]Department of Management Sciences, The University of Iowa, Iowa City, Iowa, USA
[2]College of Information Sciences and Technology, The Pennsylvania State University, University Park, Pennsylvania, USA
[3]American Cancer Society, Atlanta, Georgia, USA
[4]eBay Inc, San Jose, California, USA

Correspondence to
Dr Kang Zhao, Department of Management Sciences, The University of Iowa, S224 PBB, Iowa City, IA 52242, USA; kang-zhao@uiowa.edu

ABSTRACT

Objective Online health communities (OHCs) have become a major source of support for people with health problems. This research tries to improve our understanding of social influence and to identify influential users in OHCs. The outcome can facilitate OHC management, improve community sustainability, and eventually benefit OHC users.

Methods Through text mining and sentiment analysis of users' online interactions, the research revealed sentiment dynamics in threaded discussions. A novel metric—the number of influential responding replies— was proposed to directly measure a user's ability to affect the sentiment of others.

Results Using the dataset from a popular OHC, the research demonstrated that the proposed metric is highly effective in identifying influential users. In addition, combining the metric with other traditional measures further improves the identification of influential users.

INTRODUCTION

More than 80% of adult internet users in the USA use the internet for health-related purposes.[1] Among them, 34% read about health-related experiences or comments from others. Unlike traditional health-related websites that only allow users to retrieve information, online health communities (OHCs) allow members to share their own experience and interact with peers facing similar health problems.[2 3] Prior research has identified many benefits from OHC participations, including increased support, perceived empathy, and optimism,[4 5] as well as reduced levels of stress, depression, and psychological trauma.[6 7] Based on the Pew 2013 Health Online survey, 26% of adult internet users read or watched someone else's experience about health or medical issues in the last 12 months.[8]

The effectiveness and proper functioning of OHCs can be affected by the presence and activities of influential users (IUs),[9] who have 'the power or capacity of causing an effect in indirect or intangible ways'.[10] In general, finding IUs in online communities can help community building and management, marketing, political campaigns, and information retrieval and dissemination.[11 12] Previous research has also established that social contacts influence health-related behaviors and emotions in others.[13–15] Thus finding IUs in an OHC could have additional implications in advocating new treatments, guiding the proper use of drugs, providing necessary emotional support, and encouraging healthy lifestyles and positive attitudes.[16–19]

Classic social network theories regarding IUs focused on social network structures and have been developed along two lines: centrality metrics and influence models. Centrality metrics, such as *betweenness*,[20] *degree*,[21] *closeness*,[22] and *Pagerank*[23] quantify the importance of a node based on its structural position in a network. However, they do not directly address the extent to which one individual actually alters another's behavior, attitudes, or perceptions. Influence models characterize the dynamics of social influence using network diffusion or contagion models,[24–26] which are widely used in studying viral marketing and epidemics. IUs can be identified by finding individuals, whose infections lead to maximal diffusion.[27] Nevertheless, approaches based on influence models can be a computationally expensive undertaking, especially for dynamic online communities or networks with ever-changing topologies. Also, in these models, influence may be confounded with other factors,[28] such as homophily[29] and simultaneity.[30] Consequently, it is difficult to assess the validity of these influence models in a given context and hence difficult to assess the quality of the identified IUs.

The popularity of online communities[31] makes available detailed information on asynchronous and distributed online interactions among users. Thus two directions for identifying IUs have emerged. First, new social networks can be built to incorporate various types of inter-personal relationships beyond 'knowing', such as re-tweeting in Twitter or endorsement in product review websites. Network centrality metrics have been applied to these new networks to find IUs.[32–34] Second, individuals' behavior in online communities and their influence to others have also been studied through metrics such as log-in frequency at social networking websites,[35] bloggers' amount and styles of blogging,[36 37] and content similarity between posts in online forums.[38]

Compared to other online communities, OHCs feature unique interaction patterns because they aim to provide various types of social support. Therefore, this research looks at the impact of influence that is generated when peer support is provided, and the ways in which IUs change or alter other OHC members in the process. Our approach focuses on OHC members' emotion dynamics in online interactions. Emotions are closely related to physical health,[39] with positive emotion being considered beneficial.[40] Emotions expressed by OHC members are influenced by the emotional support they receive, as well as by the quality of information support and companionship they are given.[41] Examining members who experienced emotional changes as a result of the support they received, we identified those members whose support potentially influenced such emotional change. From this analysis, a new metric was developed to measure an individual's influence in providing social support and altering

other members' emotions in an OHC. We illustrate and evaluate the proposed method in the following sections.

METHOD

Our approach attempts to assess how online support affects OHC members' emotion dynamics in three steps: (1) it measures how inter-personal influence correlates with sentiment changes; (2) it identifies possible key contributors to sentiment dynamics in threaded discussions; and (3) it aggregates a member's possible contributions to sentiment dynamics in an OHC. We will illustrate the approach with a case study of the American Cancer Society Cancer Survivors Network (CSN). CSN is an online peer support community with over 173 000 registered members who are cancer survivors or caregivers. Virtually all site content is user-generated. The de-identified dataset used in this research is from member contributions to CSN discussion boards between July 2000 and October 2010, which is comprised of 48 779 threaded discussions with more than 468 000 posts from 27 173 unique members. Each threaded discussion starts with an *initial post*, which is published by the thread originator. This may be followed by two types of replies. Replies from other users (responders) are called *responding replies*. In many cases, the thread will contain one or more additional replies from the originator, called *self-replies*. Figure 1 shows an example of threaded discussions.

Sentiment analysis

In an OHC, user emotions cannot be observed directly, but the sentiments of their posts often reflect their emotions at the time of posting. Manually labeling the predominant sentiment expressed in every post for a large community is typically impractical, hence automatic sentiment analysis has been adopted to classify emotion embedded in texts.[42] Sentiment analysis is a text mining technique that extracts subjective information from texts, such as political opinion in tweets,[43] polarity in financial news,[44] and subjectivity in product reviews.[45]

Although there are numerous off-the-shelf sentiment analysis tools, directly using them without first checking the validity of these tools for a specific application can be problematic, because the same word related to negative sentiment in one context may not be so related in another context. For example, in the context of cancer, 'positive findings in a diagnostic test' indicates the presence of a cancer and generally expresses negative sentiment, while a 'positive news about the economic recovery' expresses positive sentiment. Thus it is desirable to use posts from an online community of interest to train sentiment classifiers, which enables the classifier to take into account context when classifying text sentiment. This research uses a classifier specifically calibrated to classify cancer forum posts into *positive* or *negative* sentiment classes. In this paper, we briefly describe how the classification model was built. More details regarding the classifier can be found in our previous research.[46]

To identify and calibrate the classification model, 298 randomly selected posts were manually labeled by two independent annotators as belonging to either the *positive* or *negative* sentiment class. Cohen's κ statistics[47] (κ=0.82) suggested high inter-annotator agreement. Then the two annotators discussed posts whose sentiment they initially disagreed on until they reached a consensus on sentiment labels. Of the 298 posts, 60% were labeled as having positive sentiment.

We extracted various lexical features (eg, the length of posts, the frequency of certain words) from the content of each post and used them as potential explanatory variables while constructing the sentiment classification model for positive and negative
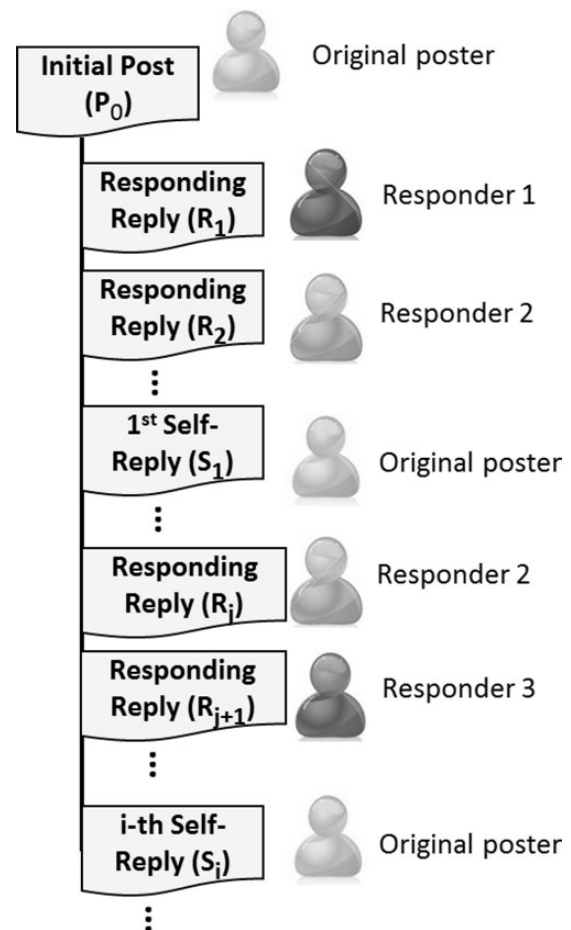

Figure 1 An example of threaded discussions with responding replies and self-replies.

sentiments (the list of features is in Section 3 of the online supplementary materials). Eight different types of classification algorithms were evaluated by using these features as inputs. AdaBoost,[48] where regression trees are used as weak learners, achieves the best performance with a classification accuracy of 79.2% (10-fold cross validation). A weak leaner is a classifier that is only slightly correlated with the true classification. AdaBoost tries to build a series of weak learners to create a strong learner that is well correlated with the true classification. Performance at this level has also been seen with other sentiment analyses of various domains, where accuracy rates ranging from 66% to 84% have been reported.[42 49] Section 4 of online supplementary materials includes more evaluation of the classifier.

The AdaBoost classification model was subsequently applied to all unlabeled posts, producing a sentiment label for each. Specifically, for each post m_i, the sentiment classification model estimates a sentiment posterior probability, $P_r(c=pos|m_i)$, which measures how likely it is that the post belongs to the *positive* sentiment class given its lexical characteristics. If $P_r(c=pos|m_i) > 0.5$, post m_i is labeled as *positive*; otherwise, it is labeled as *negative*. Figure 2 illustrates the process of sentiment classification for posts.

Sentiment dynamics

Given the assigned sentiment of all the posts, sentiment dynamics within threads were used to develop a metric that reflects each user's potential to influence others' sentiment. We focused the study of sentiment dynamics on those of thread originators,
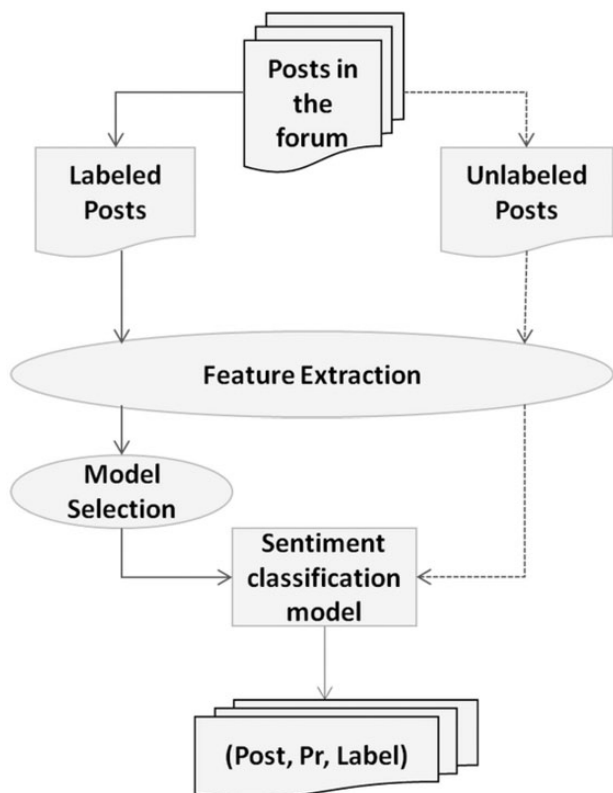
Figure 2 The process of sentiment classification for posts in the online health communities.

because they often start a thread to seek support from the community.[50] Subsequent replies from others exert some level of influence on the originator's feeling on the issue, and the sentiment of the originator's subsequent self-replies in the same thread may reflect such change.

If the thread originator does not post a self-reply to the thread she/he started, her/his change in sentiment on the issue discussed in this thread cannot be directly assessed by our method. Among the 23 000 threads, the numbers of self-replies follows the power-law distribution (figure 3). Furthermore, the distributions of positive sentiment probabilities (of initial posts) for threads without self-replies is similar to the distribution for threads with self-replies (see Section 7, online supplementary materials). Comparing the sentiment of a thread originator's initial post with her/his sentiment in subsequent self-replies could reveal the potential influence of those who respond to the thread.

The probability of thread originators' posts having positive sentiment, on average, increases within the threads they initiated (figure 4). Furthermore, the most significant change in average positive sentiment probability occurs between the initial post and the first self-reply, and the sentiment probabilities do not change much after the first self-reply in the same thread. In our analysis, the positive sentiment probability of originators' self-replies is averaged as $S_F = \sum_{i=1}^{N} p_r(pos|s_i)/N \sum_{i=1}^{N} p_r(pos|s_i)/N$, where s_i refers to one of the N self-replies from the thread originator. Similarly, the positive sentiment probability of responding replies is averaged as $S_R = \sum_{j=1}^{M} p_r(pos|r_i)/M \sum_{j=1}^{M} p_r(pos|r_i)/M$, where r_j is one of the M responding replies in the thread. The indicator of positive sentiment probability change for a thread originator is computed as $\Delta_{Pr} = S_F - S_0$, where $S_0 = p_r(c=pos|s_0)$ is the positive sentiment probability from the thread's initial post (s_0). Figure 5 plots Δ_{Pr} against S_R and demonstrates that Δ_{Pr} tends to have higher values as S_R increases (the Pearson correlation coefficient
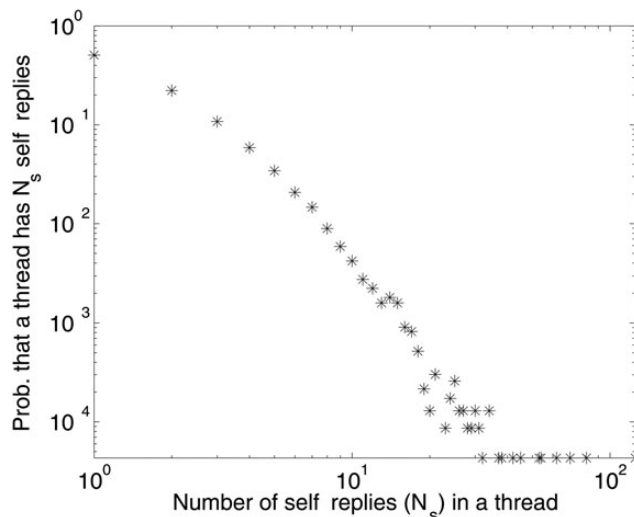

Figure 3 Log–log distribution of the number of self-replies in threads with at least one self-reply and one responding reply.

between Δ_{Pr} and S_R is $\rho=0.96$, $p \leq 0.001$). Hence, the positive sentiment probability of responding replies is positively correlated with the changes in the positive sentiment probability of the thread originator's posts.

After at least one responding reply is received, about 75% of all the thread originators who started with negative sentiment expressed a higher positive sentiment probability subsequently; among those who started with positive sentiment, 85% stayed positive. Figure 6 shows the distribution of positive sentiment probability change, Δ_{Pr}, for threads starting with negative sentiment. Only 7.9% of them have $\Delta_{Pr}<0$. The average Δ_{Pr} is 0.3811 and is significantly greater than 0 ($p<0.05$).

Influential responding replies

Having established that the sentiment of responders has an impact on changing sentiment probabilities of the thread initiator, the next issue is to use this finding to develop an approach to identify IUs. We posit that IUs post many posts that may
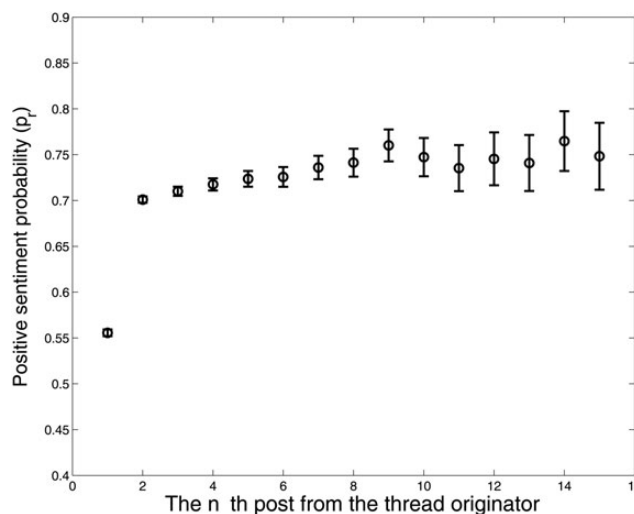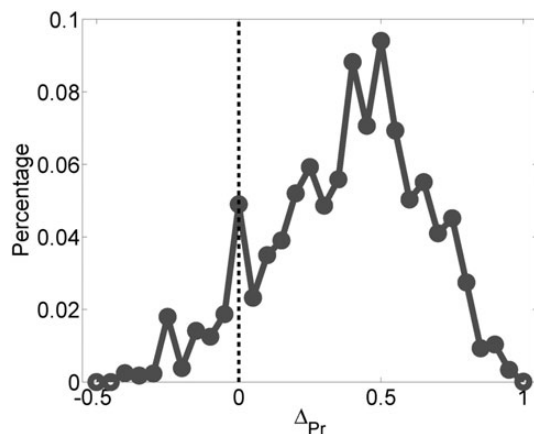

Figure 4 Sentiment change of thread originators by the number of self-replies (vertical bars are 95% CIs). The 2nd post from the originator is the 1st self-reply, hence the 2nd circle from the left-hand side denotes the average sentiment probability of 1st self-replies.
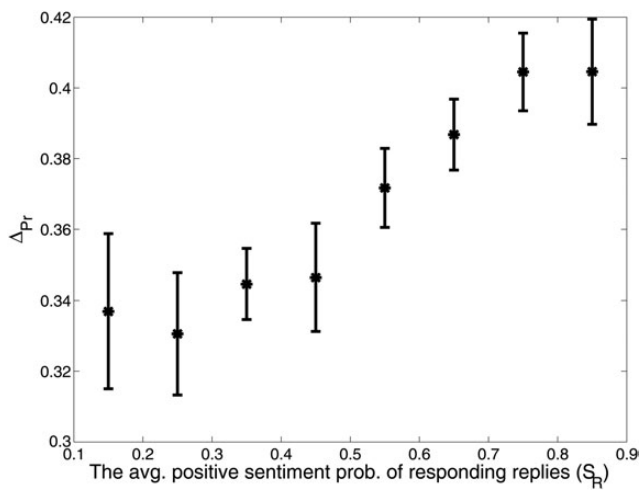
Figure 5 Change in originators' sentiment probabilities as a function of the average sentiment probabilities of responding replies (vertical bars denote 95% CIs).

affect the sentiment of thread originators and use each individual's total number of potential *influential responding replies* (IRRs) as a metric of influence.

An IRR is defined as a responding reply whose sentiment is aligned with the thread originator's change of sentiment probabilities. While all responding replies in a thread may alter sentiment probabilities of the originator's self-replies, we only consider replies that are published before an originator's first self-reply within the thread. The rationale for this is twofold. First, as discussed earlier, the average of thread originators' positive sentiment probabilities changes significantly between the initial post and the first self-reply, but they only change little afterwards (see figure 4). Thus, by examining responses in a thread before the originator posts, we are able to capture most of the sentiment dynamics from the original poster because the sentiment change happens early and additional support beyond the first self-reply does not influence the sentiment that much. Second, the temporal intervals between initial posts and first self-replies have a median value of only 17 h, with two-thirds of them below 24 h. By focusing on those early responses, we can further increase the probability that an originator's sentiment change is actually due to the responding replies and not due to support from offline friends and acquaintances, or other offline

Figure 6 The empirical distribution of Δ_{Pr} for originators who start a thread with a post that has negative sentiment.

events, such as changes in physical conditions and holiday celebrations, or simply the passage of time.

Operationally, an IRR's sentiment correlates with the thread originator's sentiment probability change between the initial post and the first self-reply. If an IRR r_j expresses positive sentiment in a thread (ie, our sentiment classifier assigns $p_r(c=pos|r_j)>0.5$ for post r_j), the originator's sentiment in the first self-reply (s_1) should be more likely to be positive compared to the initial post (s_0), that is, $p_r(c=pos|s_1)>p_r(c=pos|s_0)$; if r_j contains negative sentiment (ie, $p_r(c=pos|r_j)<0.5$), the originator's sentiment in the first self-reply should become more negative than it was in the initial post, that is, $p_r(c=pos|s_1)<p_r(c=pos|s_0)$. If there are multiple replies between the initial post and the first self-reply, we consider all IRRs to have contributed to the originator's sentiment change. This assumption is analogous to the aggregation of influence from multiple actors in threshold-based contagion models,[26][51] even though thresholds that represent individual differences on the ease of being influenced are not explicit in the definition of IRR. Formally, a responding reply r_j in the thread started by initial post p_0 is designated an IRR if and only if the following two conditions are met:

Condition 1: $T(s_0)<T(r_j)<T(s_1)$, where $T(m)$ is the publishing time of post m.

Condition 2: $p_r(c=pos|r_j)>0.5$, if $p_r(c=pos|s_1)>p_r(c=pos|s_0)$ or $p_r(c=pos|r_j)<0.5$, if $p_r(c=pos|s_1)<p_r(c=pos|s_0)$.

Note that even though IRRs are designed to assess potential influence of posts, it does not directly assess the number of posts that caused the sentiment change of thread originators. For example, a cancer survivor posting a question regarding his/her anxiety about the side-effects of a procedure can receive multiple replies with positive overall sentiments expressing encouragement and emotion support with different personal experiences regarding the side-effects. If the cancer survivor posts a reply more positive than the original post later on the thread, the change of sentiment can be caused by one, a subset, or all of the replies. IRRs do not attempt to identify such causal relationships. Also, the set of IRRs includes responding replies that express negative sentiment as well as those that express positive sentiment for two reasons. First, negative sentiment in a post is not necessarily bad if the sentiment is appropriate to the situation, such as when the originator post reports the death of a community member and replies express sympathy. A preliminary lexical search reveals that approximately 13% of initial posts contain words or expressions related to death (details are in the online supplementary materials). Second, negative sentiment may be an important factor in identifying IUs whose passion about an issue is valued by some members but antagonizes others. However, this study did not attempt to differentiate negative from positive IRRs with respect to IUs.

OHC members can be ranked based on their total numbers of IRRs. According to the metric, an OHC member with higher IRRs tends to engage in conversations with many members of the community (have more posts than other users), be prompt in replying to a new thread (before the original poster came back to the thread), and whose posts contribute to changing the sentiment probability of the thread initiator (sentimentally influential).

RESULTS

Evaluating the effectiveness of the IRR metric in identifying IUs requires an independently derived list of OHC IUs. Directly finding true IUs in a community is a challenging task that requires very good knowledge of each user's history of activities and interactions with others over an extended period of time.

For this study, the CSN community manager discussed with two CSN staff members, who monitor forum content on a full-time basis, and nominated 41 CSN community members as IUs (referred to as List-1). The list was not rank ordered because the three nominators were unable to do so reliably for so many members (when evaluating information retrieval and search engines, only relevance judgments and not totally ordered ranks are sought because of the same reason).

Although List-1 may not include all the IUs in the community, it provided a starting point for evaluating the utility of using IRR totals to identify IUs. We evaluated the ranking of users by IRR totals with Top-K recalls (also known as Recall@K). It examines how many of the nominated IUs in List-1 can be ranked within Top-K (with various K values) by a metric. If n of them are ranked within the Top-K, then the Top-K recall fraction is n/41. The higher the Top-K recall for a ranking metric, the better the performance of the metric. As a comparison, we also included many other metrics from previous literature, including individual contributions and network centralities in the post-reply network (an edge points from user B to user A if A published a responding reply in a thread started by B). Results in table 1 suggest that the performance of the IRR metric is better than that of other metrics for various values of K.

We then compared the performance of the new metric to the combined power of traditional metrics. Previous research[52] proposed several classifiers that can also be used to identify IUs (details in Section 2 of online supplementary materials). These classifiers leverage 68 features that measure users' contributions (eg, the numbers of posts and active days), network centralities (eg, degree, betweenness, and PageRank), and post content (eg, the frequency of words with positive/negative sentiment in a user's posts). These classifiers were initially built based on IU List-1 and predict whether a user outside the list is an IU. The same group of domain experts from CSN reviewed a list of 150 users who were identified by the best-performing classifier as most likely to be influential but were not included in List-1. Using criteria similar to those that generated IU List-1, they endorsed an additional 85 members as IUs (referred to as IU List-2). Performance of IRR total in identifying IUs in the combined IU List-1 and IU List-2 compared to the best-performing classifier is presented in table 2. Surprisingly, IRR total outperforms the more complicated classifier, even though the classifier combines the power of 68 other features. Furthermore, adding the IRR metric to the other 68 features resulted in the best-performing classifier: the IRR-enhanced classifier (bottom row, table 2). The IRR-enhanced classifier performs much better than the original classifier that does not use IRR.

DISCUSSIONS AND LIMITATIONS

This research develops a novel metric to identify IUs in OHCs. It focuses on sentiment change in OHC members who receive support, and assesses how the provision of that support influences OHC members' sentiment. Compared to existing metrics for IU identification, the new metric has intuitive explanations as well as good performance. Moreover, the concept of 'influential post' introduces a fundamental element of social influence at the inter-personal level. It complements the previous emphasis on analyzing 'relationship' networks with a new perspective that analyzes the conversation of actual social interactions in which influence takes place. The concept is based on the alignment of a responding reply's sentiment (inter-personal level) with the direction of the sentiment change of the thread originator at the individual level. In other words, it not only considers how much an OHC member has contributed to the community, but also how her/his contributions have influenced others.

The IRR total is also a cumulative metric for a given user, which is desirable for a dynamic online community where new content and changes in social network structure occur every day. To update a user's IRR total, we do not have to repeatedly process all historical data and construct a new social network, as when updating betweenness or PageRank centralities. Instead, we can simply find new IRRs from the user's recent posting activities since the last time her/his IRR total was calculated, and then added to the previous IRR total. Therefore, IRR total is a metric that is highly scalable for tracking IUs in very large online communities.

This research also has important implications for building a sustainable OHC that is both active and supportive. For instance, early identification of IUs in an OHC provides community managers an opportunity to publicly recognize their contributions in various ways (eg, presenting virtual badges of honor) and encourage increased participation in the OHC. This, in turn, can reinforce positive behaviors of IUs, facilitate their assumption of various community leadership roles, and thereby assure consistency of strong peer leadership. This is especially valuable when an influential community member is lost as a result of health-related factors that limit or preclude their continued involvement in the community.

As has been the case with other studies of online social influence, our approach is limited to the examination of inter-personal influence through online interactions, while influence in an OHC can occur in many other ways. For example, members who support others likely experience some improvement in their self-esteem when their support is appreciated by the intended recipient as well as other community members. OHC members may connect via other means, such as private messaging and chat rooms. They may also exchange email addresses or phone numbers and connect offline. Even though we have tried to eliminate as much offline influence as possible by focusing on the sentiment effect of prompt replies to thread originators, the sentiment change can still be due to other factors. Using observational data only, we cannot control for all variables or claim causal effects regarding original posters' sentiment change. For instance, it is possible that a reply could have caused the original poster's sentiment change without being an IRR as defined in this paper. To achieve a more comprehensive understanding of influence or to make causal inference of sentiment dynamics in large-scale OHCs, future research needs to

Table 1 Top-K recall (Recall@K) for proposed IRR total metric and other single-metric user rankings evaluated with IU List-1

Metric to rank users by	K=50	K=100	K=150
IRR total	0.512	0.732	0.829
Total number of threads initiated	0.342	0.439	0.585
Total number of posts	0.415	0.707	0.781
Out-degree in the post-reply network	0.317	0.512	0.610
In-degree in the post-reply network	0.390	0.659	0.780
Betweenness in the post-reply network	0.293	0.366	0.488
PageRank in the post-reply network	0.341	0.585	0.756
Total number of replies within 24 h after the initial post	0.487	0.707	0.781
Total number of replies before the 1st self-replies	0.434	0.634	0.756
Total number of first replies before 1st self-replies	0.487	0.707	0.781
Total number of last replies before 1st self-replies	0.439	0.609	0.756

IRR, influential responding replies; IU, influential user.

Table 2 Comparison of the performance of various approaches

	K=50		K=100		K=150	
	Recall (max=0.397)	Precision	Recall (max=0.794)	Precision	Recall	Precision (max=0.840)
Ranking by IRR total	0.349	0.880	0.627	0.790	0.762	0.640
Ranking by total replies before the 1st self-replies	0.325	0.820	0.563	0.710	0.706	0.593
The original classifier[52]	0.278	0.700	0.532	0.670	0.698	0.587
The new classifier enhanced with IRR	0.357	0.900	0.611	0.770	0.817	0.687

Evaluated with both IU List-1 and List-2. Note that there are 85+41=126 influential users. As 126>100>50, the maximum possible values for Top-50 and Top-100 recalls are not 1. Similarly, the maximum Top-100 recall is 100/126=0.794; the maximum Top-150 precision is 126/150=0.84.
IRR, influential responding replies; IU, influential user.

analyze the complex and multifaceted online interactions between OHC members at a more fine-grained level, as well as to leverage self-reported surveys or carefully designed control experiments.

As noted previously in this paper, this study does not distinguish healthy negative sentiment influences (eg, sadness due to the death of a community member) from those that are not healthy for the community (eg, members' posts that antagonize some community members or online bullying). Making such distinctions will require a more fine-grained analysis of the content of the threads and the nature of support. This can contribute to the identification of IUs who negatively impact the community and further facilitate the management of OHCs. Also, while the alignment between IRRs and original posters' sentiment dynamics is intuitive, it does not cover all the aspects of sentiment dynamics in inter-user interactions. For example, an influential reply with positive sentiment may lead a thread originator to post more about his/her negative emotion since the expression of negative emotion can be a healing process.[53] Similarly, a reply expressing empathy and sympathy may help the originator understand that she/he is not alone and contribute to better adjustment to their health condition. In addition to emotional dynamics and sentiment influence, members can influence offline behaviors such as seeking a second opinion, negotiating care, or expressing concerns with the health provider. Analyzing data from various sources and incorporating different types of influence may further improve the identification of IUs and their roles in an OHC. These are important areas for future research.

REFERENCES

1 Fox S. *The social life of health information, 2011*. Pew Research Center's Internet & American Life Project, 2011. http://pewinternet.org/~/media//Files/Reports/2011/PIP_Social_Life_of_Health_Info.pdf
2 Rolia J, Yao W, Basu S, et al. *Tell Me What I Don't Know-Making the most of Social Health Forums*. HP Labs, 2013.
3 Kumar A, Zhao K. Making sense of a healthcare forum—Smart keyword and user navigation graphs. *ICIS 2013 Proceedings*. Milan, Italy: 2013. http://aisel.aisnet.org/icis2013/proceedings/ResearchInProgress/89
4 Nambisan P. Information seeking and social support in online health communities: impact on patients' perceived empathy. *J Am Med Inform Assoc* 2011;18:298–304.
5 Rodgers S, Chen Q. Internet community group participation: psychosocial benefits for women with breast cancer. *J Comput-Mediat Commun* 2005;10. Published Online. doi:10.1111/j.1083-6101.2005.tb00268.x
6 Beaudoin CE, Tao C-C. Modeling the impact of online cancer resources on supporters of cancer patients. *New Media Soc* 2008;10:321–44.
7 Winzelberg AJ, Classen C, Alpers GW, et al. Evaluation of an internet support group for women with primary breast cancer. *Cancer* 2003;97:1164–73.
8 Fox S, Duggan M. Health Online 2013. Pew Research Center 2013. http://pewinternet.org/Reports/2013/Health-online.aspx
9 Mehra A, Dixon AL, Brass DJ, et al. The social network ties of group leaders: implications for group performance and leader reputation. *Organ Sci* 2006;17:64–79.
10 Merriam-Webster.com. influence. 2011. http://www.merriam-webster.com
11 Gladwell M. *The tipping point: how little things can make a big difference*. Back Bay Books, 2002.
12 Katz E, Lazarsfeld P. *Personal influence: the part played by people in the flow of mass communications*. Transaction Publishers, 2005.
13 Centola D. The spread of behavior in an online social network experiment. *Science* 2010;329:1194–7.
14 Christakis NA, Fowler JH. The spread of obesity in a large social network over 32 years. *N Engl J Med* 2007;357:370–9.
15 Fowler JH, Christakis NA. Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the Framingham Heart Study. *BMJ* 2008;337:1–9.
16 Anonymous. Calling all patients. *Nat Biotech* 2008;26:953.
17 Brownstein CA, Brownstein JS, Williams DS, et al. The power of social networking in medicine. *Nat Biotech* 2009;27:888–90.
18 Ma X, Chen G, Xiao J. Analysis of an online health social network. *Proceedings of the 1st ACM International Health Informatics Symposium*. ACM, 2010:297–306.
19 Valente TW. *Social networks and health: models, methods, and applications*. Oxford University Press, 2010.
20 Freeman LC. A set of measures of centrality based on betweenness. *Sociometry* 1977;40:35–41.
21 Albert R, Jeong H, Barabasi AL. Error and attack tolerance of complex networks. *Nature* 2000;406:378–82.
22 Cobb NK, Graham AL, Abrams DB. Social network structure of a large online community for smoking cessation. *Am J Public Health* 2010;100:1282–9.
23 Brin S, Page L. The anatomy of a large-scale hypertextual Web search engine. *Comput Netw ISDN Syst* 1998;30:107–17.
24 Watts D. A simple model of global cascades on random networks. *Proc Natl Acad Sci USA* 2002;99:5766–71.
25 Goldenberg J, Libai B, Muller E. Talk of the network: a complex systems look at the underlying process of word-of-mouth. *Mark Lett* 2001;12:211–23.
26 Valente TW. Network models of the diffusion of innovations. *Comput Math Organ Theory* 1996;2:163–4.
27 Kempe D, Kleinberg J, Tardos E. Maximizing the spread of influence through a social network. *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2003:137–46.
28 Aral S, Walker D. Identifying influential and susceptible members of social networks. *Science* 2012;337:337–41.
29 McPherson M, Smith-Lovin L, Cook JM. Birds of a feather: homophily in social networks. *Annu Rev Sociol* 2001;27:415–44.
30 Manski CF. Identification of endogenous social effects: the reflection problem. *Rev Econ Stud* 1993;60:531.
31 Wellman B. Net surfers don't ride alone: Virtual communities as communities. In: Smith MA, Kollock P. eds. *Communities in cyberspace*. Routledge, 1999:167–94. http://books.google.com/books?id=harO_jeoyUwC
32 Cha M, Haddadi H, Benevenuto F, et al. Measuring user influence in twitter: The million follower fallacy. Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media (ICWSM'10), 2010:10–17.

33 Zhou H, Zeng D, Zhang C. Finding leaders from opinion networks. Proceedings of IEEE International Conference on Intelligence and Security Informatics, 2009:266–8. doi:10.1109/ISI.2009.5137323

34 Zhang J, Ackerman MS, Adamic L. Expertise networks in online communities: structure and algorithms. *Proceedings of the 16th International Conference on World Wide Web*. ACM, 2007:221–30.

35 Trusov M, Bodapati AV, Bucklin RE. Determining influential users in internet social networks. *J Mark Res* 2010;47:643–58.

36 Agarwal N, Liu H, Tang L, *et al*. Identifying the influential bloggers in a community. *Proceedings of the International Conference on Web Search and Web Data Mining*. ACM, 2008:207–18.

37 Zhao K, Kumar A. Who blogs what: understanding the publishing behavior of bloggers. *World Wide Web* 2013;16:621–44.

38 Yang C, Tang X. Who made the most influence in MedHelp? *IEEE Intell Syst* 2012;27:44–50.

39 Anderson CL, Agarwal R. The digitization of healthcare: boundary risks, emotion, and consumer willingness to disclose personal health information. *Inf Syst Res* 2011;22:469–90.

40 Salovey P, Rothman AJ, Detweiler JB, *et al*. Emotional states and physical health. *Am Psychol* 2000;55:110.

41 Bambina A. *Online social support: the interplay of social networks and computer-mediated communication*. Youngstown, NY: Cambria Press, 2007.

42 Pang B, Lee L. Opinion mining and sentiment analysis. *Found Trends Inf Retr* 2008;2:1–135.

43 Mejova Y, Srinivasan P, Boynton B. GOP primary season on twitter: 'popular' political sentiment in social media. *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*. New York, NY, USA: ACM, 2013:517–26.

44 Devitt A, Ahmad K. *Sentiment polarity identification in financial news: a cohesion-based approach*. Association for Computational Linguistics, 2007:984–91. http://www.aclweb.org/anthology/P/P07/P07-1124 (accessed 3 Oct 2013).

45 Pang B, Lee L. A sentimental education: sentiment analysis using subjectivity summarization based on minimum cuts. *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*. Stroudsburg, PA, USA: Association for Computational Linguistics, 2004.

46 Qiu B, Zhao K, Mitra P, *et al*. Get online support, feel better—sentiment analysis and dynamics in an online cancer survivor community. *Proceedings of the Third IEEE Third International Confernece on Social Computing (SocialCom'11)*. 274–81.

47 Cohen J. A coefficient of agreement for nominal scales. *Educ Psychol Meas* 1960;20:37–46.

48 Freund Y, Schapire RE. A decision-theoretic generalization of on-line learning and an application to boosting. *J Comput Syst Sci* 1997;55:119–39.

49 Liu B. *Web data mining: exploring hyperlinks, contents and usage data*. Springer, 2006.

50 Joinson A. *Oxford handbook of internet psychology*. Oxford University Press, 2007.

51 Granovetter M. Threshold models of collective behavior. *Am J Sociol* 1978;83:1420–43.

52 Zhao K, Qiu B, Caragea C, *et al*. Identifying leaders in an online cancer survivor community. *Proceedings of the 21st Annual Workshop on Information Technologies and Systems (WITS'11)*. Shanghai, China, 2011:115–20.

53 Iwamitsu Y, Shimoda K, Abe H, *et al*. The relation between negative emotional suppression and emotional distress in breast cancer diagnosis and treatment. *Health Commun* 2005;18:201–15.