Engagement Capacity and Engaging Team Formation for Reach Maximization of Online Social Media Platforms

Alexander Nikolaev*
University at Buffalo
312 Bell Hall
Buffalo, New York 14260
anikolae@buffalo.edu

Shounak Gore
University at Buffalo
113 Davis Hall
Buffalo, New York 14260
shounakg@buffalo.edu

Venu Govindaraju
University at Buffalo
113 Davis Hall
Buffalo, New York 14260
venu@cubs.buffalo.edu

ABSTRACT

The challenges of assessing the "health" of online social media platforms and strategically growing them are recognized by many practitioners and researchers. For those platforms that primarily rely on user-generated content, the *reach* – the degree of participation referring to the percentage and involvement of users – is a key indicator of success. This paper lays a theoretical foundation for measuring *engagement* as a driver of reach that achieves growth via positive externality effects. The paper takes a game theoretic approach to quantifying engagement, viewing a platform's social capital as a cooperatively created value and finding a fair distribution of this value among the contributors. It introduces *engagement capacity*, a measure of the ability of users and user groups to engage peers, and formulates the Engaging Team Formation Problem (EngTFP) to identify the sets of users that "make a platform go". We show how engagement capacity can be useful in characterizing forum user behavior and in the *reach maximization* efforts. We also stress how engagement analysis differs from *influence* measurement. Computational investigations with Twitter and Health Forum data reveal the properties of engagement capacity and the utility of EngTFP.

Keywords

social networks, engagement, reach, team formation problem, influence maximization

1. INTRODUCTION

Measuring social influence online is becoming a sophisticated and granular task. Several new-age companies, e.g., Klout and PeerIndex, are in the business of evaluating the influential power of their clients' online userbases. Social media influence, defined as an individual's ability to affect the thinking or actions of peers, is the driver of success of many of today's online services, apps, ideas, etc.; the identification of influential persons, with the objective of using them to alter political preferences in a community or advertise products to increase sales, is an active research direction.

However, influence measurement is not particularly useful to an online platform manager who does not seek to exploit its userbase, but instead, strives just to maintain/encourage user activity. The purpose of most online activities including sharing pictures or posts, contributing comments or "likes" in response to peers' posts, or exchanging opinions in social forum threads is to share experiences and maintain (friendly) communication between fellow users of the network. This form of communication keeps the users together, and ultimately, defines the influential power of the platform that they use. Take as an example the downfall of MySpace and Orkut with the success of Facebook. As more and more users migrated from Orkut to Facebook, the utilization of Orkut declined, leading to its eventual abandonment. The ability to detect such turns and strategically grow online platforms is of interest to many practitioners, putting the question of assessing the platform's "health" up on the research agenda.

The recently emerging applications of social media analysis shift the focus of investigation from huge populations to well-defined, smaller subpopulations, e.g., interest-based groups, followship bases and health forums, with their own culture and purpose. Such subpopulations are of special interest to the research community, since they may develop their own mechanisms of peer pressure delivery and discourse features, and often, have significant societal value. For example, the forum analyses can aid us in drug safety signal detection efforts (similar to "early warning systems" adopted for disaster prevention [5]), disease symptom profiling, and understanding the concerns and behavior of various patient groups that communicate online [42]. Naturally, for the online health forums that prove to be effective in helping their users, there is a need for the methods and tools to attract more users.

Taking a page from RE-AIM, a widely used systematic program evaluation framework [12], we use the term "reach" to refer to an online platform's ability to attract new users, and define "engagement" as any existing user's degree of involvement in the platform's activities. The interplay between engagement and reach can be understood by leveraging on the research of the structure, stagnation, and growth of consumer markets. Studies of cascade emergence in demand-side economics have found that the customers' willingness to choose (adopt) a product grows with the number of people that have already adopted it; this increase in value, otherwise known as positive network externality, occurs with each sale of a product unit [14]. On an online platform, positive

*Corresponding author.

KDD '16, August 13-17, 2016, San Francisco, CA, USA

DOI: http://dx.doi.org/10.1145/2939672.2939681

network externalities occur in two instances: (1) when a new user joins the network and authors a new post, and (2) when an existing user contributes new user-generated content [37]. Thus, the internal growth of a platform (achieved through added user-generated content) leads to its external growth (the increase in the number of newly registered users): i.e., higher engagement leads to higher reach.

We posit that engagement occurs when a user contributes content in response to someone else's contribution(s). A network of directed relationships reflecting "which post attracts which" captures the sequence and structure of engagement (see Figure 1). An online forum grows through its users: every post (even "weather talk") fosters user "bonding" and creates a positive externality effect, even more so when it is read and responded to. The propagator – the creator of seed content – is said to engage the responder, while the responder is said to engage into the forum's activity.

Addressing the need for theory-supported quantification of engagement, we introduce the term "engagement capacity", interpreted as a user's ability to engage peers and measured as their share in the platform / forum engagement power. We view the value generated by an online social network community as a direct product of communication between its members, attributed to the submitted contributions' content and structure (order), and only in part to the volume (note that irrelevant interactions are promptly removed from pro-health sites, a majority of which are moderated).

In order to assess any individual user's engagement capacity, we employ cooperative game theory. Cooperative game theory is a branch of science that calculates "fair-share" equilibria in settings with agents that form coalitions to achieve a common goal. Its first fundamental advance is due to Shapley [32], who introduced a method for calculating fair contributed values for "players" forming unstructured coalitions, where any contributor can interact or team up with any other contributor with the objective of generating synergistic value. This is a perfect setup for the engagement ability measurement: the un-normalized and normalized (e.g., by the number of contributions or the amount of time spent online) engagement capacity can allow for fair comparison between users as they contribute to the growth of a platform.

It is important to see the potential of studying "engagement" in social networks as opposed to studying "influence" enabled by them. The latter type of analysis requires one to define how the connected agents affect each other's decisions with respect to some query; the former is query-independent. Several recent position papers stress the potential utility of engagement-focused research, citing a series of under-explored questions in this area. The beginnings of this emerging literature, along with other relevant works, are reviewed in Section 2, motivating the investigations in the present paper. Section 3 discusses the details of using cooperative game theory for engagement quantification. Section 4 introduces the concept of targeted engagement capacity and the Engaging Team Formation Problem (EngTFP). Section 5 discusses the insights shed by the engagement capacity calculations with real-world data. Section 6 gives the concluding remarks and discusses future research directions.

2. BACKGROUND AND RELATED WORK

The desire to understand the patterns in the spread of ideas, technologies, and viral product adoption in social networks motivates the modern studies of social influence mechanisms. These studies often use the term "word-of-mouth" [40] and offer diffusion-based models that explain the defining elements of innovation adoption [8], disease spread [22], and other network-driven phenomena in sociology [36] and politics [10]. The developed metrics for quantifying the ability of actors to influence their peers' decision-making, based on the structural properties of the network, in which they are embedded, are known as centrality metrics [7, 28, 23].

The work on "influence maximization" that emerged over the past decade focuses on the problem of selecting a group of the initiators (commonly termed "seeds" or "opinions leaders") to generate the largest influence cascade - by the number of adopters of a product or idea [15]. The ability to identify influential subsets of users of an online resource is particularly valuable for viral marketing [29], personalized recommendation delivery services [33], as well as microblogging [6] and health forum analyses [35].

This paper attacks a new question: "What subset of users is the most important for keeping the whole user base together?" – keeping current users engaged and keeping the platform appealing to potential new users [43].

The more recently published position papers describe the patterns of user activity that can help explain "engagement". These papers emphasize the need for quantifying engagement, and call for systematic approaches to tracking and enhancing it in calculated ways.

Lim et al. distinguish several forms of engagement in analyzing TV-viewership data [19]. They define functional engagement as real-time participation of online media platform users that results in content development; communal engagement is then presented as a consequence of functional engagement: essentially, it is what we refer to as "reach". Li et al. describe engagement as "informedness", i.e., the "ability to inform", distinguishing cognitive, relational and behavioral aspects of engagement [18]. Vydiswaran et al. discuss the reasons why users create new groups and summarize the insights for designers and researchers who desire to build better, larger communities [38]. Online social support provision is a topic of exploration of the recent papers about online health forums [39]; indeed, such forums are known to successfully deal with behavioral aspects of health: they help "e-patients" maintain diets, cope with side effects of medications, quit smoking and remain sober [43].

Several prescriptive research ideas stem from engagement analyses. Stearns et al. call for "smart thread" recommender systems to boost the reach of online health forums [34]. The logic of Bernabe-Moreno et al. [5] suggests that the engagement measurement can inform "early warning" systems signaling changes in the health of an online platform. Another practically useful idea lies in maintaining and/or growing the userbase of a platform by identifying and nurturing the users that are most responsible for its (high) engagement.

The questions of measuring the motivation of a user to contribute to a forum thread first came up in the studies of Question-and-Answer forums [3]. The simple measures, e.g., the number of upvotes on websites like StackOverflow, or the number of "likes" on Facebook and Twitter, reflect forum utilization and signal user interest. Natural language processing (NLP) works with the text of a post to understand its success in attracting responses [4, 39]. In an effort pursuing a similar objective to ours, Lehmann et al. use Twitter data to study "transient crowd-forming" and find most *in-*

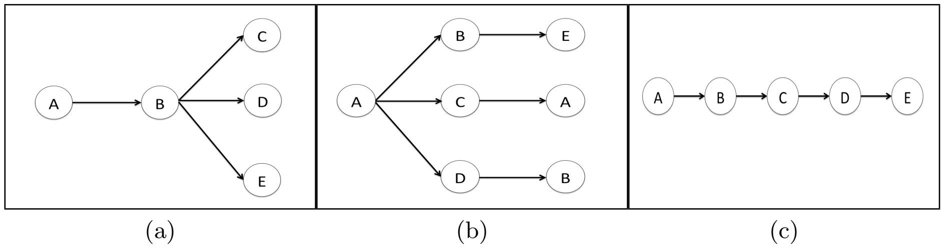

Figure 1: Examples of communication threads of different structure.

fluential/engaging curators of news propagation based solely on text analysis: they find common topics of interest which bind the transient crowd together [16, 17]. The NLP methods, however, do not yet account for the dynamics of users creating posts and interacting with each other; they require much human supervision, and tend to be slow in processing big data. Overall, the above measures and methods serve the purposes different from quantifying how much merit a user deserves for engaging peers in conversation.

The studies of engagement, accounting for the dynamics of interaction, have emerged recently. The terms "engagingness" and "responsiveneness" are used in modeling email message chains [26, 25]. In a study of Twitter, the count of re-tweets is taken as a measure of engagingness [1]. These works, however, are based on direct responses from one user to another, with indirect communication (engaging one user through another) not accounted for. Zhao et al. explore how the structure of online interactions affects the sentiment of the messages posted by a given user [44], considering the situations where the same user appears at least twice in the same thread, e.g., asks a question and then replies to the respondent(s). The work of Yang and Tang on understanding the mechanisms of online influence [43] goes further in exploring the reasons why a user may succeed in attracting responses to their posts. They recognize – as do we – that the contributions of forum thread posts to making the thread engaging must depend on the posts' positions within the thread, with the immediate predecessor deserving more credit for attracting a new (follow-up) post in the thread.

This paper introduces new metrics to quantify engagement by fully exploiting the dynamics of user interactions in online communication-based platforms. The following section describes our methodological approach in detail.

3. MEASURING THE ENGAGEMENT CAPACITY OF A USER

The seminal work on calculating fair shares in cooperatively created value for agents forming unstructured coalitions is due to Shapley [32]. His work was extended by Myerson [21], and later, by other researchers, to problems with constraints on cooperation structure and frequency [27, 24, 41, 9].

The most relevant to our work are the extensions that recognize that in certain situations, the value generated by a coalition may depend on the order in which players interact [24, 30]. We extend and adapt this line of work to render it useful in thread-based communication context, by allowing for multiple interactions between users (as players) in forum threads, accounting for contribution (forum post) order.

Consider a setting where forum users cooperatively generate engagement value: when they contribute posts that attract more posts, they promote the forum growth and increase its reach achieved through positive externality effects. Each post contributed in response to other post(s) is assumed to generate one unit of engagement value for the forum. We define the engagement capacity of a user as their share in the forum's total engagement value, based on all the threads the user has contributed to. The engagement capacity of a user can be positive if and only if the user has been active in at least one thread; the passive readers (lurkers) do not add engagement value to the forum. Moreover, the users who contribute, but are never responded to, have zero engagement capacity; indeed, if a forum consisted of only such members, it would generate no communication at all.

Based on a single thread, the engagement capacity of a user depends on the position(s) of the user's contribution(s) in the thread's flow. As such, a thread starter is (in part) responsible for engaging all the users that have replied to the thread: a user whose post generates further posts gets partial engagement value credit from all such contributions. On the other hand, the posts immediately preceding a newly contributed post get more credit for engaging it than the posts appearing earlier in the thread: otherwise, the new contribution would be expected to have occured earlier or be directly responding to those older posts.

Let us look closely at the examples in Figure 1: assume that three forum threads have been created by the same five users. User A originates each thread. User B replies to A's original messages in all the three threads. In the first thread (Figure 1a), users C, D and E reply to B, instead of directly replying to A. In the second thread (Figure 1b), users C and D also reply to A's question directly. The third thread (Figure 1c) consists of a single "line" thread branch. Based on the directed networks in Figure 1, the averages of some standard centrality metrics are reported in Figure 2 (the detailed explanation of how the engagement capacity metric is computed in the figure and the discussion of how the metrics relate to each other will be given in the subsequent sections). A methodological approach to computing engagement capacity is presented next, followed by a discussion of some computational aspects of the engagement analysis.

3.1 Cooperative Games on k–Coalitions

Nowak and Radzik [24] and Sánchez et al. [30] explain that for a transferable utility (cooperative) game in a directed graph, the worth of a coalition can depend not only on the individual properties of coalition members but also on the order in which they interact in a coalition. Nowak and Radzik [24] define the value Ψ^{NR} that generalizes the Shapley value for transferable utility games. For a game with player set N and value function v, where player subsets $S \subset N$ form ordered coalitions $T \equiv (i_1, ..., i_{|T|})$ from the set of all possible ordered coalitions $\pi(S)$, this value for player i amounts to

$$\Psi_i^{NR}(N, v) = \sum_{S \subset N \setminus \{i\}} \sum_{T \in \pi(S)} \frac{(|N| - |T| - 1)!}{|N|!} (v(T \cup i) - v(T)),$$

which, with $\Omega(N)$ as the set of all possible ordered subsets of N, $H(T)$ as the set of players in coalition T, and $i(T)$ as the position of player i in coalition T, can be concisely re-written as

$$\Psi_i^{NR}(N, v) = \sum_{T \in \Omega(N), i \in H(T), i(T) = |T|} \frac{\Delta_v^*(T)}{|T|!},$$

where $\Delta_v^*(T)_{T \in \Omega(N), T \neq \emptyset}$ are termed the generalized coefficients of v, also known as the coordinates of v in the generalized unanimity basis [13]. Sanchez and Bergantinos [30] define another Shapley value extension, distinguishing coalitions by the players' positions *within* them,

$$\Psi_i^{SB}(N, v) = \sum_{T \in \Omega(N), i \in H(T)} \frac{\Delta_v^*(T)}{|T| \, (|T|!)}.$$

More recently, del Pozo et al. [9] defined a generalization of Ψ^{NR} and Ψ^{SB} as a parametric family of functions $\{\Psi^{\alpha}\}_{\alpha \in [0,1]}$, where the value generated by an ordered coalition is shared proportionally to the positions of its members,

$$\Psi_i^{\alpha}(N, v) = \sum_{T \in \Omega(N), i \in H(T)} \Delta_v^*(T) \frac{\alpha^{|T| - i(T)}}{|T|! \sum_{j=0}^{|T|-1} \alpha^j}.$$

We extend and adapt the line of work on ordered transferable utility games [24, 30, 9] to render it useful in the thread-based communication context, where interactions take place between online platform users (as players) as they contribute posts to forum threads in response to each other, in sequence, and possibly, multiple times to the same thread. To this end, define k-coalition as a connected ordered sequence of player *appearances*. As with the coalitions used to define $\{\Psi^{\alpha}\}_{\alpha \in [0,1]}$, a k-coalition is distinguished not only by its membership, but also, by its ordering. Similarly to $i(T)$, we let $i(T, k)$, $k = 1, 2, ..., K$, denote the position of the k-th appearance of player $i \in H(T)$ in k-coalition $T \in \Omega^K(N)$, with $\Omega^K(N)$ denoting the set of all k-coalitions in which any given player can appear at most K times. The value generated by a k-coalition is shared proportionally to the positions of the appearances of the coalition members,

$$\Psi_i^{K\text{-}\alpha}(N, v) = \sum_{T \in \Omega^K(N), i \in H(T), k=1,..,K} \Delta_v^*(T) \frac{\alpha^{|T| - i(T,k)}}{|T|! \sum_{j=0}^{|T|-1} \alpha^j}. \tag{1}$$

The family of parametric functions $\{\Psi^{K\text{-}\alpha}\}_{K \in I^+, \, \alpha \in [0,1]}$ encompasses, as special cases, the conventional Shapley value as well as the functions Ψ^{NR}, Ψ^{SB}, and $\{\Psi^{\alpha}\}_{\alpha \in [0,1]}$.

The concept of engagement capacity is now introduced, in conjunction with the term *engaging subthread* in forum communication. Define subthread as an uninterrupted chain (sequence) of posts of a thread that begins with the first post of the thread; in a directed tree graph, representing a forum thread, every path that begins at the root is a subthread. A subthread is called *engaging* if it is succeeded by at least one post in its respective thread. The number of forum posts contributed in response to or following up on another post (or a sequence of posts) over the whole forum constitutes a total engagement value generated by the forum users; the share of each user in this value is called engagement capacity: it measures each user's ability to engage peers, computed retrospectively, i.e., based on their past activity records.

Given an online forum, let N denote the set of all the users in the forum's userbase. Define (N, v, P) as a game on the set of all the forum's subthreads P. A subthread-restricted game (N, v_U, p) can be interpreted as the game where users $U \subset N$ contribute posts to form $p \in P$; this setup is similar, but not exactly the same, to a game in a communication "situation" described in [9]. Given the forum's historical data, set K to be the largest number of posts contributed by the same user to any subthread; let the *value function* $\Delta_v^*(T)$ return a total number of posts immediately succeeding such *engaging* subthreads $p \in P$ that have the same membership, size and structure as k-coalition $T \in \Omega^K(N)$. The engagement capacity of forum user $i \in N$ is the value that (1) returns for this user as a solution of the game (N, v, P) and is hereafter denoted by η_i, or by $\eta_{i,F}$ to specify that the computation is based on a restricted set of forum threads F.

Note that the coefficient $\alpha \in [0, 1]$ in $\{\Psi^{K\text{-}\alpha}\}_{K \in I^+, \, \alpha \in [0,1]}$ captures the engagement share tradeoff between thread contributors, as they manage to attract each new post to "their" thread. As such, if a new post is submitted in response to multiple (preceding) consecutive posts, then its immediate predecessor gets more credit for attracting it, with the credit to the earlier predecessors discounted by the factors of α, α^2, α^3, etc., respectively. Note that this property of engagement capacity complies with an observation that the posts, newly attracted to a (long) forum thread, tend to respond to the latest prior contributions in this thread [43].

3.2 Calculating Engagement Capacity

Consider the example in Figure 1a, and denote the thread depicted on it by "(a)". Using (1), the engagement capacity of user A is found to be

$$\eta_{A,(a)} = 1 + 3 * \frac{\alpha}{\alpha + 1},$$

as user A gets the engagement value of 1 by contributing to engaging subthread A with $\Delta_v^*(\{A\}) = 1$, and gets $\frac{3\,\alpha}{\alpha+1}$ by contributing to engaging subthread AB with $\Delta_v^*(\{AB\}) = 3$; meanwhile, subthreads ABC, ABD and ABE are not engaging. The engagement capacity of user B is

$$\eta_{B,(a)} = 3 * \frac{1}{1 + \alpha}.$$

Note that users C, D and E have zero engagement capacity in this example; also, the total engagement value generated and shared is equal to four, i.e., to the number of posts contributed by all the users in response to their peers. Table 1 gives the engagement capacity values for each thread in Figure 1, with $\alpha = 1$, i.e., giving all the subthread propagators equal credit for engaging a responder. In these examples, the early contributors have higher engagement capacities than the late ones. However, in general, this is not necessarily the case; for example, with $\alpha = 0$, $\eta_{A,(a)} = 1$ and $\eta_{B,(a)} = 3$.

The engagement capacity values can be interpreted directly, or upon normalization. One sensible way is to normalize by the number of contributed posts to identify the

User	(a)	(b)	(c)
A	2.5	4.5	2.083
B	1.5	0.5	1.083
C	0	0.5	0.583
D	0	0.5	0.25
E	0	0	0

Table 1: Engagement capacity values for the threads in Figure 1.

users whose contribution content is engaging irrespective of the volume.

The case where a forum post is immediately followed (in the same thread) by another post of the same user is worth a separate discussion. The treatment of such occurrences is up to the researcher: e.g., one may choose to "merge" such posts and treat them as one. However, in a moderated forum, one may assume that every contribution is distinct, i.e., makes a new point; since every post creates positive externalities, rewarding the user (with a share in the total engagement value) for both posts makes sense. Another issue that becomes apparent in the considered example is that a user may gain a high engagement capacity by engaging themselves (as opposed to others): this issue will be addressed in Section 4, with the introduction of "targeted engagement capacity."

Engagement capacities of users can be, rather conveniently, dynamically computed and updated, as new content is added, without the need to redo the analysis for the whole history of the forum/platform every time it changes. Recall that each new user post submitted in response to another post, or sequence of posts, brings in one unit of engagement value to the communication thread and to the platform as a whole. A newly added post increases the engagement capacity values of all the users contributing to the subthread leading to the new post (but not including it). Consider a k-coalition $T' \in \Omega^K(N)$ with the same membership, size and structure as this subthread. The increase in the engagement capacity of user $i \in H(T')$, resulting from the addition of the new post, is given by

$$\Delta_i = \sum_{k=1,2,\ldots,K} \frac{\alpha^{|T'|-i(T',k)}}{\sum_{j=0}^{|T'|-1} \alpha^j}. \tag{2}$$

Equation (2) specifies how every new contribution changes the engagement capacity values of the forum users. This equation can be used in real time to efficiently track the contributed engagement dynamics. Note that the expression in (2) can be evaluated in $O(n)$ time: its denominator requires finding the subthread succeeded by a new post, and its numerator requires the information about the positions of user appearances in this subthread.

4. TARGETED ENGAGEMENT CAPACITY AND ENGAGING TEAM FORMATION

Equation (2) of Section 3 specifies how to fairly distribute a unit engagement value, brought to the platform with any new post submitted into an existing thread, between the thread contributors. Naturally, some users may be succesfull in engaging certain peers and not so successful in engaging others. This realization highlights the value of solving a game of engaging one given user j: the peers of j would split the engagement value generated by the j's responses to their posts or post sequences. In this case, only those subthreads that are followed up on by the posts of j would be considered engaging.

More generally, one can address the question of evaluating the ability of a given set of users, $V \in N$, to engage another set of users, $W \in N$, in forum communication. To this end, we introduce the term *targeted engagement capacity*: denoted by $\eta_{V\to W}$, it is defined as the sum of the shares allocated to the members of V in the game of engaging the members of W. In the setting of the game on k-coalitions, described in Section 3.1, one has

$$\eta_{V\to W} = \sum_{i\in V, j\in W} \sum_{i\in H(T), k=1,..,K} \Delta_v^*(T) \frac{\alpha^{|T|-i(T,k)}}{|T|! \sum_{j=0}^{|T|-1} \alpha^j}, \tag{3}$$

where $\Delta_v^*(T)$ returns a total number of posts by the members of W immediately succeeding such *engaging* subthreads $p \in P$ that have the same membership, size and structure as k-coalition $T \in \Omega^K(N)$. Note that generally, $\eta_i \neq \eta_{i\to N\setminus i}$: the difference between these quantities is $\eta_{i\to i}$ – it indicates how much user i tends to engage in back-and-forth conversations as opposed to occasionally contributing to multiple forum threads. Note also, that targeted engagement capacity can be updated dynamically, as well as the originally defined engagement capacity, by a trivial extension of (2).

Table 2 shows the targeted engagement capacity values, per equation (3), for the users contributing to the threads in Figure 1. The columns of Table 2 are for the the users whose engagement capacity is calculated and the rows are for the targeted users. The sum over each column gives a total engagement capacity of the respective user over all the three thread. Note that $\eta_{A\to A} > 0$: user A self-engaged in the second thread.

User	A	B	C	D	E
A	0.5	0.00	0.5	0.00	0.00
B	3.5	0.00	0.00	0.5	0.00
C	2.00	1.00	0.00	0.00	0.00
D	1.833	0.833	0.333	0.00	0.00
E	1.25	1.25	0.25	0.25	0.00

Table 2: Targeted engagement for the threads in Figure 1

The ability to measure how successful a particular user is in engaging other particular users allows us to attack the following question: "What group (team) of active users is most engaging?" This question is of special interest to any growing online platform, since such teams of users can be rewarded, encouraged, and assisted in further increasing peer engagement and retention. To help this cause, the Engaging Team Formation Problem (EngTFP) is introduced.

The EngTFP objective is to select a subset of users with a maximal targeted engagement capacity towards all the other users: $\max_U \eta_{U\to N\setminus U}$. This problem can be formulated with additional constraints, e.g., those specifying which historical engagement data to take into account and/or the maximum size of the subset to be selected (e.g., $|U| \leq b$).

To set up an instance of EngTFP, first, use (3) to compute all the pairwise targeted engagement capacity values, $\{\eta_{i\to j}\}_{i\in N, j\in N, i\neq j}$. Note that, in general, $\eta_{i\to j} \neq \eta_{j\to i}$ for $i \neq j$. Let X_i, $i \in N$, be binary decision variables such that for any $i \in N$, $X_i = 1$ if i is selected into U, and zero otherwise. Let Y_{ij}, $i \in N$, $j \in N$, $i \neq j$, be auxiliary binary variables that would be set to 1 if and only if $X_i = 1$ and

$X_j = 0$ (at the same time). Then, EngTFP is given,

$$\begin{aligned} \max \quad & \sum_{i \in N} \sum_{j \in N, j \neq i} \eta_{i \to j} Y_{ij} \\ \text{s.t.} \quad & Y_{ij} \leq X_i, \quad \forall\, i, j, \\ & Y_{ij} \leq 1 - X_j, \quad \forall\, i, j, \\ & X_i, Y_{ij} \in \{0, 1\}, \quad \forall\, i, j. \end{aligned}$$

A non-linear formulation of EngTFP does not use any auxiliary variables and maximizes $\sum_{i \in N, j \in N, j \neq i} \eta_{i \to j} (X_i^2 - X_i X_j)$; this formulation is quadratic but not convex, and hence, does not allow for a dominant convex decomposition [20], confirming that EngTFP is combinatorially challenging.

EngTFP is a special case of MAX2SAT problem, known to be NP-Hard; for a review of the algorithmic work on MAX2SAT, see [11]. The complexity of EngTFP is due to the fact that, once user i is selected into a team, the team members can no longer be rewarded for engaging i. Indeed, if a platform decides to pay some users for helping grow it, the users to be paid should be good at attracting other users but not each other.

Again, we underline the difference between the EngTFP and the influence maximization problem. Influence maximization strives to enhance a certain effect (e.g., change in political opinion, product adoption, etc.) in a pre-existing and known user network. Meanwhile, EngTFP solutions aim at helping a platform maintain or build its network or non-network userbase. The EngTFP advises a decision-maker on what users should be rewarded, virtually (e.g., via badges, titles, points) or physically (e.g., via gift cards, discount codes, cash), for igniting communication, which the users achieve through content generation, question asking, social support provision, information exchange, etc.

5. COMPUTATIONAL INVESTIGATIONS

This section reports the engagement capacity analyses conducted with the data from two active online communication platforms differing in purpose. We begin with an account of the data collection, and then, present the numerical findings.

5.1 Data Collection

We collected the forum contribution records from an online healthcare platform, one of the biggest active and freely accessible online sites for pro-health social networking. It has about 200 social support forums and about as many "ask an expert" forums. The website has close to 3 Million active and inactive threads and attracts about 8 Million visitors every month. The users in these health forums (HF) interact through discussion boards, contribute personal journal entries exploiting weight and mood tracker features, and post notes on their friends' home pages. The HF data most relevant to the present study are those of the users' interactions on moderated discussion forums.

A web crawler was used to collect the data of the contributed posts, and their sequences, in the "anxiety", "cholesterol control" and "weight-loss & dieting" forums. The anxiety forum has about 33500 threads dating back to the early 2000s. The cholesterol control forum has about 300 threads. The weight-loss & dieting forum has about 7000 threads. Each thread consists of at least one post, however, not all the threads have replies; the unanswered threads do not affect the engagement capacity computation. After cleaning up threads with no replies we end up with a database of about 48800 unique users for the HF data. Note that some of these users have been active on more than one forum.

A data tree is built for each analyzed thread. In each such tree, every user contribution is represented by a node; a directed link is drawn from node A to node B, if B replies to node A. This tree captures the relations between the users who interact in a particular thread discussion. The tree structure also tells us who initiates a particular discussion (e.g., by posting a question), who engages the maximum number of people by furthering a discussion with comments and who interacts frequently in a particular conversation. The engagement capacities computed and presented hereafter are based on both forums (all their threads) together.

Another dataset was created using Twitter, based on about 20,000 tweets, with their re-tweets, related to the 2014 FIFA World Cup. A tweet thread consists of a tweet and all of its re-tweets. A directed link is drawn from node A to node B, if user B re-tweets user A's tweet. Thus, we have a number of communication threads where every root is a particular original tweet and the other nodes are re-tweets. This tree structure, like the HF trees, allows us to identify an initiator, the most engaging user and the person who interacts most frequently in any particular conversation. A total of 31,467 unique Twitter users contributed to these threads.

5.2 Calculations for the Example in Figure 1

Using the dynamic approach to engagement capacity measurement with the collected data, we first assess how the different influence-related measures behave, as compared to the proposed scheme of measuring user engagement.

Figure 2 reports the values of the different metrics for the example in Figure 1, considering that these three threads constitute a whole forum. In Figure 2, the users are arranged over the horizontal axis in the descending order of their engagement capacity values. The other metrics were evaluated in Python [31]. The experiments were run on a Macbook Pro machine with an Intel-i5 2.3GHz processor.

As expected, the engagement capacity values depend on the number of users engaged, the length of the communication branches, and the frequency of communication. The deeper a branch goes, the more value each of the upstream users gains. User A, involved in both engaging a lot of users (Figure 1a) and seeding deeper communication branches (Figure 1c), enjoys the highest engagement capacity. User B also has a high engagement capacity value as it engages a fair number of users directly (Figure 1b) and also gets partial credit whenever the depth of the branch increases (Figure 1c). User E has zero engagement capacity due to failing to engage any peers – this user always stops the communication.

The results in Figure 2 suggest that the calculated engagement capacity values are in line with the intuition. Engagement capacity positively correlates with out-degree: indeed, high engagement capacity signals a high ability of a user to spread information, and thus, highly engaging users are expected to be connected to more peers. Engagement capacity also correlates with betweenness centrality, since the latter indicates how capable users are of transferring information between otherwise disconnected subgroups. The page rank and in-degree behave very differently than engagement capacity since they focus on describing the information flows into a user, not the other way around.

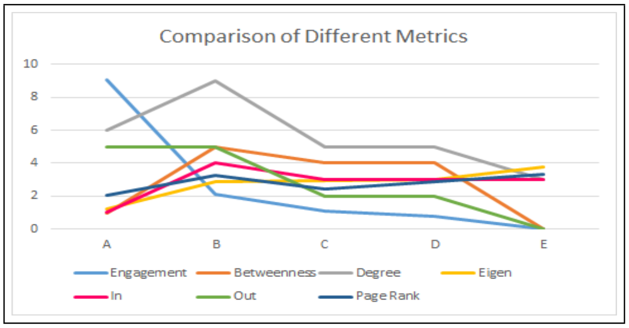

Figure 2: Engagement capacity and network-based metric values for the users of a forum with three threads as depicted in Figure 1 (best viewed in color).

	coeff	std err	t	P
betweenness	3.7676	0.166	22.695	<0.001
degree	3.1465	0.121	26.068	<0.001
eigen	4.2928	0.169	25.352	<0.001
page_rank	10.4100	0.395	26.405	<0.001

Table 3: Linear regression outputs for engagement capacity, with centrality metrics as predictors, based on HF data.

	coeff	std err	t	P
betweenness	0.6716	0.010	68.658	<0.001
degree	0.9212	0.013	71.803	<0.001
eigen	1.0924	0.016	68.685	<0.001
page_rank	1.5579	0.020	78.509	<0.001

Table 4: Linear regression outputs for engagement capacity, with centrality metrics as predictors, based on Twitter data.

5.3 Calculations for the HF and Twitter data

We now turn to the collected HF and Twitter data. In all the subsequent analyses, the (targeted) engagement capacities were computed with $\alpha = 0.99$.

Figure 3(a) shows that the distribution of engagement capacity is Gaussian in the HF and Twitter user bases. Such a good visual fit is surprising, and suggests that this metric may be capturing/revealing something about the personality of each user: indeed, personal characteristics/abilities/talents, e.g., IQ, are typically normally distributed in humans. The engagement capacity values of Twitter users are generally higher than those of the HF users because of the higher overall level of communication (i.e., retweets contributed vs. posts submitted). Figure 3(b) shows that the distribution of the number of contributions per user on each platform is likely a power law: this is common in social media analyses.

Figure 4 depicts the engagement capacity values, normalized by the number of contributions (posts), for the users in the anxiety forum (Figure 4(a)), cholesterol forum (Figure 4(b)), weight loss & dieting forum (Figure 4(c)), and Twitter data (Figure 4(d)). The anxiety forum has the largest number of contributions, however, the average engagement value per contribution in this forum is noticeably lower than in the other forums. The large number of contributions in this forum is attributed to the many discussions it hosts; however, the thread originators tend to leave the threads once they get their concerns resolved (some of them continue socializing but typically with small groups of peers). This is in line with the purpose and function of many social support forums, where users are more comfortable sharing their experiences/thoughts one-on-one, going back and forth in a conversation, so the threads get long. Consequently, the engagement values gained by these users, per contribution, are lower compared to other forums or Twitter, where a single message can engage a larger number of readers/followers. The cholesterol forum exhibits an opposite dynamics: here, the average engagement per contribution is very high, while the number of contributions per user is much smaller. This suggests that the users tend to post content which is found useful by multiple peers, attracting many responses from different sources. One reason for this behavior could be that cholesterol-related issues are much more generic (and hence, one good post can answer the concerns of many), compared to the anxiety-related issues where each case may be unique, i.e., patient-specific, and requires individual attention. The weight loss & dieting forum has a smaller variance for the distribution of engagement values per contribution.

In summary, the discussed figures verify that the user engagement capacity values do not entirely depend on the number of posts; moreover, the correlation between these two quantities is positive on Twitter but negative on HF, signaling that social support provision is difficult to maintain just by increasing the activity. Also, there is a larger variance in the engagement per contribution in the Twitter data, which can be attributed to the shorter and broader communication trees, compared to HF trees. The figures also show that some users manage to be very engaging even though they contribute to relatively few communication trees.

Figures 5(a)-(d) show the average self-engagement against the number of contributions. Users are said to self-engage when they contribute to same threads multiple times, reacting to the peers' responses to their own responses. For example, in Figure 1(b), user C replies to user A, and then, user A replies back to user C. Thus, user A self-engages and increases their own engagement capacity. Figure 5(a) shows that such situations happen frequently in the anxiety forum, thus giving the users higher self-engagement values. The phenomenon of self-engagement is consistent with and reflects the nature/type of social support provision.

Next, the various metrics used for measuring influence are compared against the presented engagement capacity metric, based on the Twitter data: see Figure 6. The Twitter users are first arranged in the increasing order of their engagement capacity values. Then, they are partitioned by the bins of equal width, separated by the percentile points (forming a total of 100 bins). The horizontal axis in Figure 6 contains the percentile values. The highest value in each bin (for each particular metric) was used to plot the graphs in Figure 6. The vertical axis shows the metric values. Figures 6(a)-6(e) present the engagement capacity, page rank, eigenvalue centrality, betweenness centrality and degree centrality values, respectively, for the Twitter users. The results with the HF data are similar, and omitted for brevity. Per Figure 6, none of the other metrics can replace engagement capacity. In order to reveal any relationship between engagement capacity and the other metrics, we perform a simple linear regression and observe a high R^2 value – around 87% for each of the data sets. This shows that engagement capacity captures multiple aspects of what it means for a node to be important in a directed communication tree, as judged by the metrics traditionally used to measure influence. Tables 3 and 4 report the found regression coefficients, standard errors, t-statistic values and p-values for the regression analyses with the HF and Twitter data, respectively.

Now we identify the top 10 engaging users for both HF

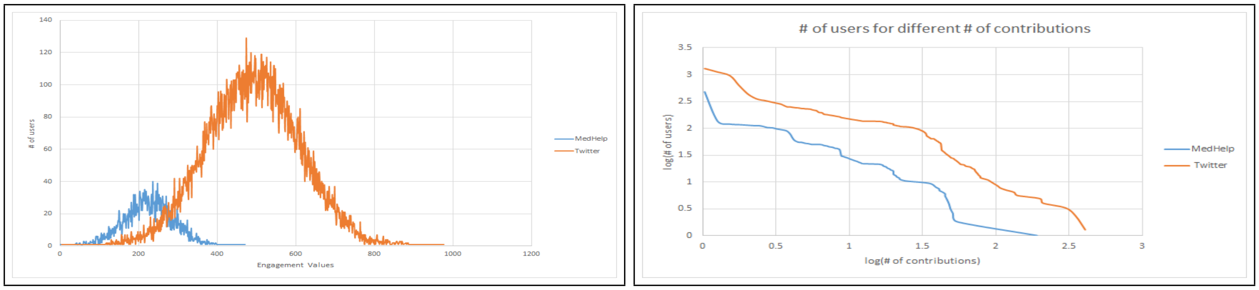

Figure 3: Distribution of engagement capacity values for HF and Twitter users (best viewed in color).

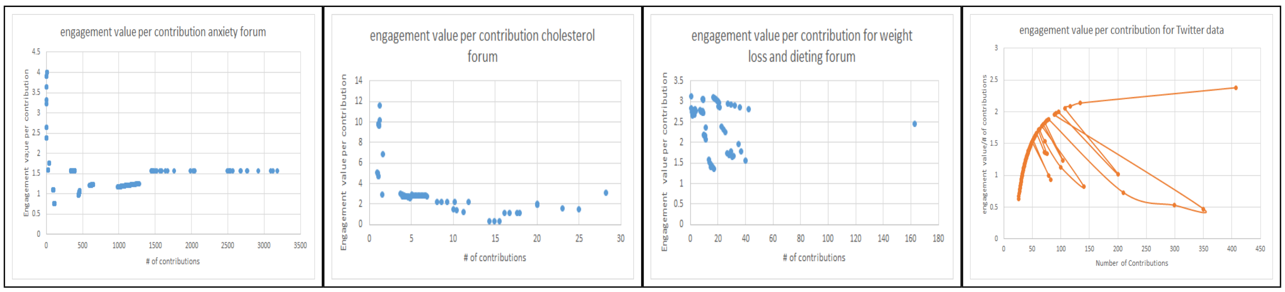

(a) (b) (c) (d)

Figure 4: Engagement capacity per user contribution, for the users ordered by the increasing contribution volume.

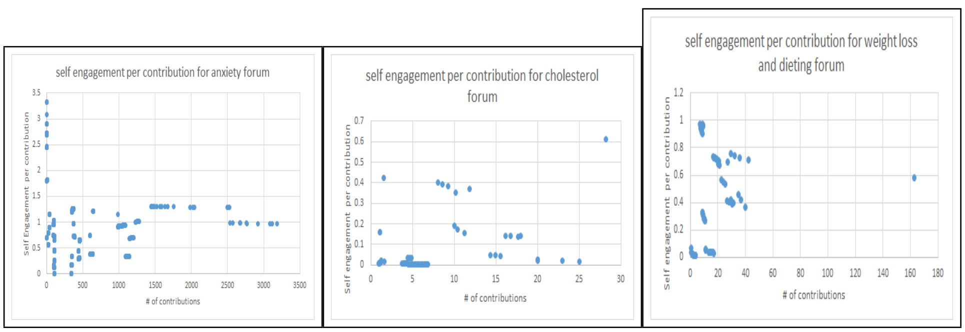

(a) (b) (c)

Figure 5: Self engagement capacity per user contribution, for the users ordered by the increasing contribution volume.

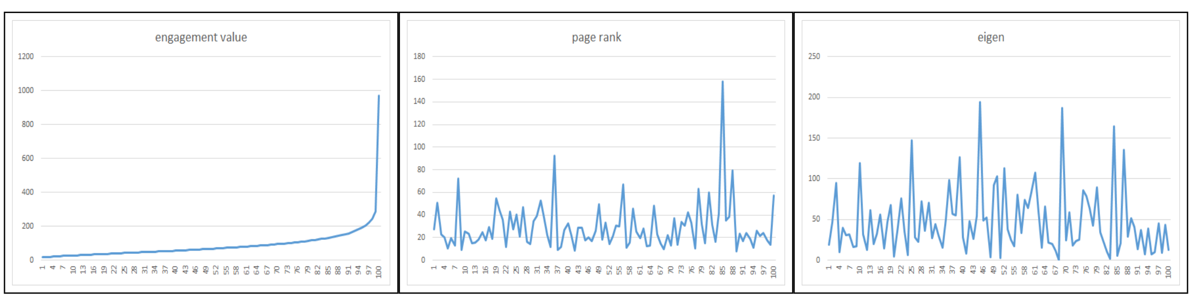

(a) (b) (c)

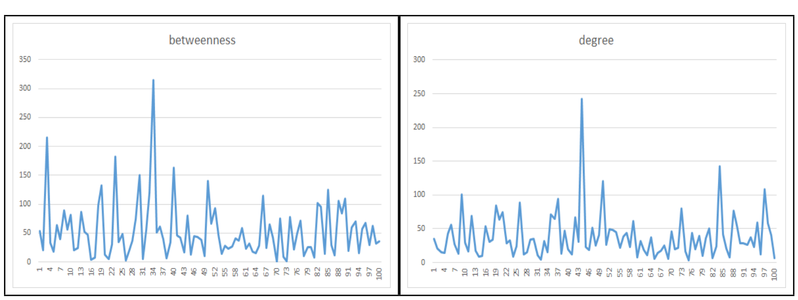

(d) (e)

Figure 6: Comparison of different metrics for the Twitter users arranged according to the increasing engagement capacity.

	A	B	C	D	E	F	G	H
A	3.9	2.3	4.3	4.5	1.2	3.8	1.3	1.3
B	2.9	2.8	3.5	1.9	0.9	1.4	2.5	3.4
C	1.9	3.8	4.6	0	0	1.7	3.4	0
D	2.5	4.4	5	0	2.4	2.6	0.6	2.3
E	1.4	4.4	1.1	1.7	1.5	4.7	2.1	0.2
F	1.2	1	2.2	1	4	5	0.6	0.5
G	0	2.4	3	4.6	4.7	2.2	2.4	0.9
H	2	0	1.5	0.8	2.7	4.5	3.5	1.6

Table 5: Targeted engagement values for top 8 engaging HF users

k	# of users in top 10	Time
2	2	$\sim$ 8 hours
3	2	$\sim$ 9.5 hours
4	2	$\sim$ 10 hours
5	3	$\sim$ 11 hours
unconstrained	5	$\sim$ 25 hours

Table 6: EngTFP results for the best team of size k.

and Twitter. The HF users can upvote each other's posts; an upvoted post receives a special label. Depending on the number of times their posts are upvoted, a user can be recognized as a top-contributor by a forum. Four of the top 10 engaging users we identify turn out to be such top contributors. This shows that knowledgeable users may earn a high engagement value, albeit without a guarantee. We also find the top 10 tweet creators in the Twitter dataset. Two of these were the accounts associated with news agencies. For an event like the FIFA World Cup, a news agency seems like the right place where other users would go for information. Three of the remaining eight users have a large number of followers and re-tweets. Engagement capacity thus helps reveal the engaging users both in terms of the content they serve as well as the followers they attract.

5.4 Targeted Engagement Capacity Values

Next, we explore what the targeted engagement capacity measure can capture, with the calculations for the top eight engaging users. Using equation (3), we summarize the results for the HF data in Table 5. In this table, users A through H are the top eight most engaging users in each forum, with A being the most engaging, B the second most engaging, etc. The columns in this table are for the users whose engagement is calculated and the rows are for the targeted users. Note that $\eta_{i \to j} \neq \eta_{j \to i}$ for $i \neq j$, i.e., the targeted engagement value from one user i to another user j is not the same as that from user j to user i. Some of the targeted engagement values are zero, signaling no interaction between the user whose engagement capacity was evaluated and the targeted user – in this order. For example, in Table 5, the value for row C and column D is zero. This means there was no communication in the HF forums where D replied to C directly or indirectly. The results with the Twitter data are similar, and omitted for brevity.

5.5 Experiments with EngTFP

Finally, we present the results of solving the EngTFP applied to the HF data. A total of 1,000 random users were selected from the HF dataset, with the objective of finding the most engaging subset of size k among them. The results for several EngTFP instances, solved as integer programs in the SCIP optimization suite [2], are summarized in Table 6. In the instances, the value of k was varied between two and five, and in one instance, k was unrestricted. Comparing the best teams' members against the list of top 10 most engaging HF users, it can be seen the EngTFP does select some highly engaging users but not all. This is because by selecting a pair of users that manage to engage each other, one loses much of targeted engagement. Interestingly, the unconstrained problem optimizes at k as low as 14. Out of the 14 users selected, five are among the top 10 most engaging users, while four others have low engagement values, which indicates that the EngTFP selects some obscure users in order to maximize the reach. The last five users of the 14 have mid-ranged engagement values. Table 6 also shows the time SCIP took to solve each instance. It should be noted that the reported figures included the time to compute the pairwise targeted engagement capacity values for every user pair, before formulating and solving each instance.

6. CONCLUSION

This paper introduces engagement capacity – a metric that serves to measure the ability of online platform users to engage each other in communication on the platform. To this end, we employ co-operative game theory. The presented dynamic method of engagement capacity calculation eliminates the need to re-calculate it from scratch with the addition of new threads and posts to a forum.

The reported experimental results show that engagement capacity can reveal well-interpretable facts about the nature of online communication on different platform types. The observations that one can obtain by calculating the self-engagement and the engagement capacity values, normalized by the number of contributions, are especially insightful. The engagement capacity distribution in a userbase reveals the different dynamics of communication and engagement in two social media, differing in purposes, in HF and Twitter.

We show that targeted engagement capacity can help evaluate the ability of one user to engage another. The targeted engagement studies may facilitate the research of the mechanisms of engagement. Finally, we show how one can use the EngTFP to identify the users who are critical to a platform's success. This work is practically valuable: many online services can rely on it to reward their users in calculated ways.

The paper provides a foundation for future research into *why* engagement occurs. With the engagement capacity metric at hand, various measurable parameters can be explored as its predictors shedding light into user behavior. One might be able to classify engaging users and predict (targeted) engagement capacity by studying its dynamics over time, and by taking into account the activity levels of the users and their followers, their attributes, network positions, behavior, features of posted texts (e.g., sentiment, emotions), etc. For example, studies of engagement predictors (explanatory variables) could distinguish the impact of news media accounts on engagement from that of regular uses; one could determine whether the media texts are any more creative than those of the regular users. The goal of such efforts would be to recognize the *innate abilities* of users to ignite communication, and encourage further activity of such users.

Further experimental work is in order to study how the fluctuations in forum engagement value and user engagement capacity over time correlate with ensuing fluctuations in new user arrival rates, existing user activity levels and turnover. It would also be interesting to see whether these dynamics differ on the Q&A versus social support platforms.

Finally, work can be expanded on more efficient implementations of various versions of constrained EngTFP, coupled with deeper theoretical investigations into its properties. The membership in the most engaging user group, per EngTFP, is likely to vary over time for a given platform. It would be worthwhile to study these dynamics more closely.

7. REFERENCES

[1] P. Achananuparp et al. Who is retweeting the tweeters? modeling, originating, and promoting behaviors in the twitter network. *ACM TMIS*, 3(3):13, 2012.

[2] T. Achterberg. Scip: Solving constraint integer programs. *Mathematical Programming Computation*, 1(1):1–41, 2009.

[3] L. A. Adamic et al. Knowledge sharing and yahoo answers: everyone knows something. In *17th WWW*, pages 665–674. ACM, 2008.

[4] E. Agichtein et al. Modeling information-seeker satisfaction in community question answering. *ACM TKDD*, 3(2):10, 2009.

[5] J. Bernabé-Moreno et al. Caresome: A system to enrich marketing customers acquisition and retention campaigns using social media information. *Knowledge-Based Systems*, 80:163–179, 2015.

[6] F. Bonchi et al. The meme ranking problem: Maximizing microblogging virality. *Journal of Intelligent Information Systems*, page 29, 2013.

[7] S. P. Borgatti. Identifying sets of key players in a social network. *Computational & Mathematical Organization Theory*, 12(1):21–34, 2006.

[8] J. S. Coleman et al. *Medical innovation: A diffusion study*. Bobbs-Merrill Indianapolis, 1966.

[9] M. del Pozo et al. Centrality in directed social networks. a game theoretic approach. *Social Networks*, 33(3):191–200, 2011.

[10] F. DeroıÌĹan. Formation of social networks and diffusion of innovations. *Research Policy*, 31(5):835–846, 2002.

[11] M. Fürer et al. *Exact Max 2-Sat: Easier and Faster*. Springer, 2007.

[12] R. E. Glasgow, , et al. Evaluating the public health impact of health promotion interventions: the re-aim framework. *American journal of public health*, 89(9):1322–1327, 1999.

[13] J. C. Harsanyi. A simplified bargaining model for the n-person cooperative game. *International Economic Review*, 4(2):194–220, 1963.

[14] M. L. Katz et al. Network externalities, competition, and compatibility. *The American economic review*, pages 424–440, 1985.

[15] D. Kempe et al. Maximizing the spread of influence through a social network. In *9th ACM SIGKDD*, pages 137–146. ACM, 2003.

[16] J. Lehmann, C. Castillo, M. Lalmas, and E. Zuckerman. Finding news curators in twitter. In *Proceedings of the 22nd international conference on World Wide Web companion*, pages 863–870. International World Wide Web Conferences Steering Committee, 2013.

[17] J. Lehmann, C. Castillo, M. Lalmas, and E. Zuckerman. Transient news crowds in social media. 2013.

[18] T. Li et al. Corporate twitter channels: The impact of engagement and informedness on corporate reputation. *International Journal of Electronic Commerce*, 18(2):97–126, 2013.

[19] J. S. Lim et al. How social media engagement leads to sports channel loyalty: Mediating roles of social presence and channel commitment. *Computers in Human Behavior*, 46:158–167, 2015.

[20] C. C. Moallemi et al. Convergence of min-sum message-passing for convex optimization. *IEEE Transactions on Information Theory*, 56(4):2041–2050, 2010.

[21] R. B. Myerson. Graphs and cooperation in games. *Mathematics of operations research*, 2(3):225–229, 1977.

[22] M. E. Newman. Spread of epidemic disease on networks. *Physical Review E*, 66(1):016128, 2002.

[23] A. Nikolaev et al. On efficient use of entropy centrality for social network analysis and community detection. *Social Networks*, 40:154–162, 2015.

[24] A. S. Nowak et al. The shapley value for n-person games in generalized characteristic function form. *Games and Economic Behavior*, 6(1):150–161, 1994.

[25] B.-W. On et al. Mining interaction behaviors for email reply order prediction. In *ASONAM 2010*, pages 306–310. IEEE, 2010.

[26] B.-W. On et al. *Engagingness and Responsiveness Behavior Models on the Enron Email Network and Its Application to Email Reply Order Prediction*. Springer, 2013.

[27] G. Owen. Values of games with a priori unions. In *Mathematical economics and game theory*, pages 76–88. Springer, 1977.

[28] L. Page et al. The pagerank citation ranking: bringing order to the web. 1999.

[29] M. Samadi et al. A subjective evidence model for influence maximization in social networks. *Omega*, 2015.

[30] E. Sanchez, , et al. On values for generalized characteristic functions. *Operations-Research-Spektrum*, 19(3):229–234, 1997.

[31] D. A. Schult et al. Exploring network structure, dynamics, and function using networkx. In *7th Python in Science Conferences (SciPy 2008)*, volume 2008, pages 11–16, 2008.

[32] L. S. Shapley. A value for n-person games. Technical report, DTIC Document, 1952.

[33] X. Song, , et al. Information flow modeling based on diffusion rate for prediction and ranking. In *16th WWW*, pages 191–200. ACM, 2007.

[34] M. Stearns et al. Towards evaluating and enhancing the reach of online health forums for smoking cessation. *Network Modeling Analysis in Health Informatics and Bioinformatics*, 3(1):1–11, 2014.

[35] X. Tang et al. Identifing influential users in an online healthcare social network. In *ISI 2010*, pages 43–48. IEEE, 2010.

[36] T. W. Valente et al. Network models of the diffusion of innovations. *Nursing Times*, 90(35):52–3, 1994.

[37] T. van Mierlo et al. Mapping power law distributions in digital health social networks: Methods, interpretations, and practical implications. *Journal of medical Internet research*, 17(6), 2015.

[38] V. V. Vydiswaran et al. User-created groups in health forums: What makes them special. In *Proc Conf Weblogs and Social Media (ICWSM) Assoc Adv Artif Intell*, pages 515–24, 2014.

[39] X. Wang et al. Social support and user engagement in online health communities. In *Smart Health*, pages 97–110. Springer, 2014.

[40] W. H. Whyte Jr. The web of word of mouth. *Fortune*, 50(1954):140–143, 1954.

[41] E. Winter. The shapley value. *Handbook of game theory with economic applications*, 3:2025–2054, 2002.

[42] C. C. Yang. Patient centered healthcare informatics.

[43] C. C. Yang et al. Estimating user influence in the medhelp social network. *IEEE Intelligent Systems*, 5(27):44–50, 2012.

[44] K. Zhao et al. Finding influential users of online health communities: a new metric based on sentiment influence. *Journal of the American Medical Informatics Association*, 21(e2):e212–e218, 2014.