Trust And Distrust Across Coalitions - Shapley Value Centrality Measures For Signed Networks

Varun Gangal
IIT Madras
vgtomahawk@gmail.com

Abhishek Narwekar
IIT Madras
abhisheknkar@gmail.com

Balaraman Ravindran
IIT Madras
ravi@cse.iitm.ac.in

Ramasuri Narayanam
IBM Research, India
ramasurn@in.ibm.com

Abstract

Signed social networks are social networks with edges indicative of both trust and distrust. We propose novel game theoretic centrality measures for signed networks, which first generalize degree-based centrality from nodes to sets, and then compute individual node centralities using the concept of Shapley Value. We derive closed form expressions for the Shapley Value for most of these measures. Moreover, we demonstrate that some of these measures give improved AP (average precision) compared to net positive in-degree for the task of detecting troll users, in the Slashdot signed network.

1 Introduction

Signed networks [1] are networks with both positive and negative interactions between nodes. They may occur as explicit edges, as in the Slashdot [2] and Epinions [3] social networks, or be inferred from interactions such as administrator elections on Wikipedia [4] or conversations on Twitter [5]. Mathematically, a signed network can be specified as (V, E^+, E^-), where V is the set of vertices, with E^+ and E^- being sets of directed edges of the form (a, b), denoting *a trusts b* or *a distrusts b*, for E^+ and E^- respectively.

In a social network, a centrality measure assigns each node a value, which denotes its importance within the network. The notion of importance may vary based on the application, leading to a wide variety of such measures, based on position [6], betweenness [7, 8] and prestige[9]. A centrality measure for a signed network needs to incorporate two sources of information and the interplay between them - namely the positive and negative edges. This, in addition to the imbalance in real-world signed networks between positive and negative edges, makes defining signed network centrality measures a non-trivial task.

A simple centrality measure for signed social networks, first proposed in [2], is the simple net positive in- degree, also called Fans Minus Freaks (FMF) centrality measure - where fans are nodes with positive edges to pointing to the node under consideration, while freaks are the one with negative edges. Other measures, such as eigenvector centrality [9] and PageRank [10], have also been generalized to signed networks.

A disadvantage of some of these centrality measures is that they consider every node in isolation when computing the centrality. This ignores the synergy between nodes, where a node is important by virtue of its combination with other groups of nodes. Moreover, ignoring the synergy can also make some of these centrality measures vulnerable to attacks wherein groups of nodes work together, as noted in [11], to boost their individual centralities or reduce other nodes's centralities. One approach to incorporate this synergy is to define a cooperative game, which assigns a value to every possible subset of nodes $C \subseteq V$ given by the characteristic function $\nu(C)$ of the game. The value assigned to a node is a weighted sum of the marginal contributions it makes to the values of all possible subsets, also known as the Shapley Value (SV). SV based centrality provides an intuitive way of capturing a node's centrality in combination with different groups of other nodes in the network.

1.1 Shapley Value Centrality - Preliminaries

A cooperative game is defined by a set of agents $A = \{a_1, a_2 \ldots a_N\}$ and a characteristic function $\nu() : P(A) \mapsto R$, where $P(A)$ represents the power set of A, and R is the set of real numbers. $\nu(C)$ essentially maps every $C \subseteq A$ to a real number, which represents the payoff of the subset, or the coalition [1]. Generally, $\nu(C)$ is defined such that $\nu(\phi) = 0$, where ϕ is the null set, which corresponds to the coalition with no agents.

Shapley value, first proposed in [12], is a way of distributing the payoff of the grand coalition (the coalition of all agents) amongst each agent. It proposes that the individual payoff of an agent should be determined by considering the marginal contribution of the agent to every possible coalition it is a part of. Such a scheme of distribution is also found to obey certain desirable criteria.

Let $\pi \in \Pi(A)$ be a permutation of A, and let $C_\pi(i)$ denote the coalition of all the predecessors of agent a_i in π, then the Shapley value is defined as

$$SV(a_i) = \frac{1}{|A|!} \sum_{\pi \in \Pi(A)} (\nu(C_\pi(i) \cup a_i) - \nu(C_\pi(i)))$$

For defining a SV based centrality measure on a graph $G = (V, E)$, we consider $A = V$. $\nu(C)$ is defined such that it represents the a measure of power/centrality/influence of the coalition C. The centrality of each node v_i is then given by its Shapley value $SV(v_i)$.

1.2 Computing the SV

Earlier methods for SV based centrality, such as [13] used a Monte-Carlo sampling based approximation to compute $SV(v_i)$ for each v_i. This essentially involves uniformly sampling a large number of permutations from $\Pi(V)$, and then finding the average marginal contribution of v_i as per the definition above. This approach is expensive (since one has to sample a large number of permutations), as well as inexact. Moreover, since the number of permutations grows as $O(n!)$, where $n = |V|$ is the number of vertices in the graph.

[14] first proposed that by defining $\nu(C)$ conveniently, one could compute the SV exactly in closed form. This approach is naturally more preferable than the MC sampling based one, although it requires defining ν, such that SV can be easily derived using arguments from probability and combinatorics.

2 Contributions

To the best of our knowledge, ours is the first work to define cooperative game theoretic centrality measures for signed social networks. Moreover, we are also the first to evaluate such measures for a centrality-based ranking task, in traditional or signed networks. Earlier works have evaluated these measures for other tasks such as influence maximization [13], or selecting gatekeeper nodes [15]. Here, we consider the task of ranking users to detect trolls or malicious users in a social network. Intuitively, these users are expected to have a a highly negative reputation amongst the users of the network. In the availability of ground truth about trolls, one can evaluate a signed network centrality measure by considering how low these "trolls" rank in a list of users ranked according to the measure.

3 Measures

We now define several Shapley value based centrality measures for directed, signed networks, given by $G = (V, E^+, E^-)$, where each edge $(a, b) \in E^+$ denotes a *trusts* b and each edge $(a, b) \in E^-$ denotes a *distrusts* b. We take G to be a signed, directed network, since in the context of Slashdot, a

[1]The terms subset and coalition are used here interchangeably

positive/negative edge denotes that user a approves/disapproves or trusts/distrusts user b's content. Hence, the directionality is of importance.

Each of our measures is based on a different definition of $\nu(C)$. We denote the positive and negative in-degrees by $d^+_{in}(V)$ and $d^-_{in}(V)$, and the corresponding out-degrees by $d^+_{out}(V)$ and $d^-_{out}(V)$. The positive and negative in-neighbor sets are denoted by $N^+_{in}(V)$ and $N^-_{in}(V)$, and the corresponding out-neighbor sets by $N^+_{out}(V)$ and $N^-_{out}(V)$.

3.1 Game Definitions & Closed Form Expressions

3.1.1 Fans Minus Freaks (FMF)

This is a simple generalization of the degree centrality measure to directed, signed networks. More formally,

$$FMF(v_i) = d^+_{in}(v_i) - d^-_{in}(v_i)$$

We attempt to generalize this measure to sets of nodes by appropriate definitions of $\nu(C)$, and then compute individual centralities of nodes using the SV of ν.

3.1.2 Net Positive Fringe (NPF)

We first define the sets $\nu^+(C)$ and $\nu^-(C)$ [2]. $\nu^+(C)$ is the set of all nodes v_j such that

- v_j has atleast one positive out-neighbor $v_i \in C$ OR
- $v_j \in C$

The second clause results from the intuitive assumption that v_i always trusts itself. $\nu^-(C)$ is the set of all nodes v_j such that v_j has atleast one negative out-neighbor $v_i \in C$. Note that v_j itself may be present inside or outside the coalition, and this does not affect $\nu^-(C)$. The characteristic function $\nu(C)$ is given by

$$\nu(C) = |\nu^+(C)| - |\nu^-(C)|$$

We can think of the $\nu(C)$ as the difference between two characteristic functions, namely $|\nu^+(C)|$ and $|\nu^-(C)|$. Hence the $SV(v_i)$ of ν will be the difference between the Shapley values $SV^+(v_i)$ and $SV^-(v_i)$ for the characteristic functions $|\nu^+(C)|$ and $|\nu^-(C)|$. Let us refer to these games as Game 1 and Game 2 respectively.

- **Game 1**
  Consider a permutation π sampled uniformly from $\Pi(V)$. Let C be the set of nodes preceding v_i in the coalition. Let v_j by any node $\in V$. Note that v_j could also be v_i itself. We denote B_{v_i,v_j} as the random variable denoting the contribution made by v_i through v_j. Here, "v_i through v_j" means that as a result of v_i being added to C, what is the effect on the membership of v_j in $\nu^+(C)$. The Shapley Value of $SV(v_i)$ of v_i, is then given by $\sum_{v_j \in V} E[B_{v_i,v_j}]$. Note that since we are taking the expectation over a permutation drawn uniformly from $\Pi(V)$, it is equivalent to averaging over every possible permutation.

  Now, we can easily see that B_{v_i,v_j} can be non-zero only if $v_j \in v_i \cup N^+_{in}(v_i)$, since in all other cases, v_j can neither be added or removed from $\nu^+(C)$ as a result of v_i being added. Now, consider the case where $v_j \in v_i \cup N^+_{in}(v_i)$. $E[B_{v_i,v_j}]$ will be equal to the fraction of permutations in which v_i is able to add v_j to the set $\nu^+(C)$. This can happen only if v_j is neither itself in C, nor is any other out-neighbor of v_j. Note that as a result, the event of v_i adding v_j only depends on the ordering of these $d^+_{out}(v_j)+1$ nodes within the permutation. Hence , we can directly consider the ordering of these nodes, ignoring the other nodes in our calculation. v_i must be the first node amongst these $d^+_{out}(v_j)+1$ in the permutation π, for it to contribute v_j. Hence,

$$\begin{aligned} E[B_{v_i,v_j}] &= \frac{(d^+_{out}(v_j))!}{(d^+_{out}(v_j)+1)!} \\ &= \frac{1}{(d^+_{out}(v_j)+1)} \end{aligned}$$

[2] Note the slight overloading of ν here. ν^+ and ν^- denote sets, not real values

Now, $SV^+(v_i)$, the Shapley value of Game 1, is consequently given by

$$SV^+(v_i) = \sum_{v_j \in v_i \cup N_{in}^+(v_i)} \frac{1}{(d_{out}^+(v_j)+1)}$$

- **Game 2**
  Using arguments similar to Game 1, we get

$$SV^-(v_i) = \sum_{v_j \in N_{in}^-(v_i)} \frac{1}{(d_{out}^-(v_j))}$$

  Note that the +1 term in the denominator for the $SV^+(v_i)$ expression is not present here, since the node v_j cannot add itself to $\nu^-(C)$

The final expression of $SV(v_i)$ for the NPF game is thus given by

$$SV(v_i) = \sum_{v_j \in v_i \cup N_{in}^+(v_i)} \frac{1}{(d_{out}^+(v_j)+1)} - \sum_{v_j \in \cup N_{in}^-(v_i)} \frac{1}{(d_{out}^-(v_j))}$$

The time taken to compute the NPF, T_{NPF} would be given by

$$T_{NPF} = O(\sum_{v \in V}(1 + d_{in}^+(v) + d_{in}^-(v)))$$
$$T_{NPF} = O(V + E)$$

3.1.3 Fringe Of Absolute Trust (FAT)

Note that we refer to $\nu^+(C)$ and $\nu^-(C)$, as defined in the previous game. In this game, a node is included in the coalition's value if it satisfies both the conditions below

- It either belongs to the coalition or has a positive out neighbor in the coalition.
- It does not have a negative out neighbor in the coalition.

The intuition underlying this measure is that every node contributing to the set's value should be such that it does not distrust any node in the set. Most signed networks have more positive edges than negative ones. For instance, Slashdot has only 23.9% of its edges marked as negative. Hence, the negative edges may be interpreted strongly as an explicit "vote of distrust". Note that distrusting even a single member in the coalition removes a node from the coalition's value.

Previous work proposes similarly motivated solutions to overcome imbalance. For instance, in [16], the authors undersample the positive edges to be equal to the negative edges.

If node $v_j \in N_{in}^+(v_i) \cup v_i$, then B_{v_i,v_j} is +1 if v_i is the first of any of the out-neighbours of v_j (positive or negative) or the node v_j itself to occur in the permutation. This is because if a $v_k \in N_{out}^+(v_j) \cup v_j$ is in C, without any negative out neighbor of v_j being in C, then v_j already is such that $v_j \in \nu^+(C) - \nu^-(C)$. Also, if any negative out neighbor of $v_j \in C$, then v_j can never belong to $\nu^+(C) - \nu^-(C)$, hence adding v_i to C would have no effect. This argument holds good even if $v_i = v_j$. Therefore, for $v_j \in N_{in}^+(v_i) \cup v_i$

$$E[B_{v_i,v_j}] = \frac{1}{d_{out}^+(v_j) + d_{out}^-(v_j) + 1}$$

Now consider $v_j \in N_{in}^-(v_i)$. We can see that $B_{v_i,v_j} \neq 0$ iff

- A node $v_k \in N_{out}^+(v_j) \cup v_j$ belongs to C
- No negative out-neighbor v_k of v_j belongs to C.

In fact, $B_{v_i,v_j} = -1$ if both the conditions above are satisfied. The expectation is given by

$$E[B_{v_i,v_j}] = -\frac{\sum_{x=1}^{x=d_{out}^{+}(v_j)+1} \binom{d_{out}^{+}(v_j)+1}{x} x!(d_{out}^{total}(v_j) - x)!}{(d_{out}^{total}(v_j) + 1)!}$$

where $d_{out}^{total}(v_j) = d_{out}^{+}(v_j) + d_{out}^{-}(v_j)$

Note that $\binom{n}{r}$ represents the number of ways of choosing r distinct things from n distinct things. Since computing factorials for large values can become infeasible in code due to limits on the value of variables, we simplify the expression into a product of fractions form.

$$E[B_{v_i,v_j}] = \frac{-1}{d_{out}^{total}(v_j) + 1} \sum_{x=1}^{z=d_{out}^{+}(v_j)+1} \frac{\prod_{\alpha=1}^{\alpha=x} (d_{out}^{+}(v_j) - x + \alpha)}{\prod_{\alpha=1}^{\alpha=x} (d_{out}^{total}(v_j) - x + \alpha)}$$

The final expression for $SV(v_i)$ for NPF is given by

$$SV(v_i) = \sum_{v_j \in v_i \cup N_{in}^{+}(v_i) \cup N_{in}^{+}(v_i)} B_{v_i,v_j}$$

The complexity of computing FAT, T_{FAT} would be

$$T_{FAT} = O(\sum_{v \in V} d_{in}^{+}(v_i) + 1 + d_{in}^{-}(v_i)(\Delta_{out}^{+})^2)$$

$$T_{FAT} = O(V + E + E(\Delta_{out}^{+})^2))$$

where Δ_{out}^{+} is the maximum positive out degree.

3.1.4 Negated Fringe Of Absolute Distrust (NFADT)

This game is in some sense like FAT, but with the roles of distrust and trust reversed. $\nu(C)$ here is given by $-|\nu^{-}(C) - \nu^{+}(C)|$. The negative sign is because $|\nu^{-}(C) - \nu^{+}(C)|$ would be a measure of disrepute (negative reputation). We omit the expression here for the sake of brevity. The complexity expression of NFADT would be similar to that of T_{FAT}, with Δ_{out}^{+} replaced by Δ_{out}^{-}.

3.1.5 Net Trust Votes (NTV)

The intuition underlying this measure is that the collective importance of a group of nodes is the net number of "votes" or edges in its favour, by nodes outside the group.

Given a coalition C, let E^{+} be the set of positive in-edges from a node outside the coalition to a node in the coalition. Similarly, E^{-} is the set of negative in-edges from a node outside the coalition into the coalition. Note that we do not consider internal edges in either term. In the NTV game, $\nu(C)$ is given by $|E^{+}| - |E^{-}|$. Let us now consider the derivation of a closed form expression for this game. Consider the node v_i being added to the C. v_i can contribute to the value of $|E^{+}| - |E^{-}|$ in four different ways, as stated below

1. Positive out-edges from v_i to nodes $v_k \in C$. These edges become internal when v_i is added to the coalition, decreasing the value of $|E^{+}| - |E^{-}|$ by 1.
2. Negative out-edges from v_i to nodes $v_k \in C$. These edges become internal when v_i is added to the coalition, increasing the value of $|E^{+}| - |E^{-}|$ by 1.
3. Positive in-edges from $v_k, v_k \notin C$ to v_i. These edges become a part of E^{+}, increasing $\nu(C)$ by 1
4. Negative in-edges from $v_k, v_k \notin C$ to v_i. These edges become a part of E^{-}, decreasing $\nu(C)$ by 1.

Consider case 1. This case only happens for v_i, $v_j \in N_{out}^{+}(v_i)$. For this event to happen, $v_j \in C$. In other words, it should precede v_i in the permutation. This will happen in exactly half the permutations, and will result in a contribution of $\frac{-1}{2}$. The cumulative contribution as a result of case 1 will be $\sum_{v_j \in N^+_{out}(v_i)} -\frac{1}{2} = -\frac{d^+_{out}(v_i)}{2}$.

Symmetrically, in case 2, the cumulative contribution will be $\sum_{v_j \in N^-_{out}(v_i)} \frac{1}{2} = \frac{d^-_{out}(v_i)}{2}$.

Case 3 can only happen for v_i, $v_j \in N^+_{in}(v_i)$. Here v_j should follow v_i in the permutation i.e. it should not be in C. This will happen in exactly half the permutations. Hence, the cumulative contribution will be $\sum_{v_j \in N^+_{in}(v_i)} \frac{1}{2} = \frac{d^+_{in}(v_i)}{2}$. Symetrically in case 4, we have the cumulative contribution given by $\sum_{v_j \in N^-_{in}(v_i)} -\frac{1}{2} = -\frac{d^-_{in}(v_i)}{2}$. Summing over the contributions from the 4 cases we get

$$SV(v_i) = \frac{1}{2}(d^+_{in}(v_j) - d^-_{in}(v_j)) - \frac{1}{2}(d^+_{out}(v_j) - d^-_{out}(v_j))$$

This expression also exhibits a relation to status theory [1], where a positive edge from a to b, indicates that b has a higher status than a, while a negative edge indicates the opposite. Thus for a node a, there are $d^+_{in}(v_j) + d^-_{out}(v_j)$ nodes with a status lower than it, and $d^-_{in}(v_j) + d^+_{out}(v_j)$ nodes with a status higher than it. The expression above, is in some sense, an indicator of the node's status relative to its neighbors. Since this expression involves only the node degrees, $T_{NTV} = O(V)$, provided we maintain the in-degrees and out-degrees separately.

3.2 k-Hop NPF

We can generalize the **NPF** measure to k hops, considering the sets of in-neighbours $N^+_{in}(v_i, k)$ and $N^-_{in}(v_i, k)$ to be the set of neighbors within a distance of k-hops using in-edges. The notion of a path being positive or negative is determined using the principle "The enemy of my enemy is my friend", motivated by balance theory [1]. However, unlike the traditional balance theory setting, we consider the direction of the edges i.e. we only consider paths where edges are pointing from the source to the destination.

The sign of the path is given by the product of its edge signs. Note that $N^+_{in}(v_i, k)$ and $N^-_{in}(v_i, k)$ may now have a non-zero intersection unlike the simple one-hop neighbor sets. Hence, we do not generalize the FAT and $NFADT$ measures to the k hop case, since these measures do not consider the neighbor sets independently.

3.3 Strong Balance Shapley Value (SBSV)

This measure has its motivation in balance theory [17], which has also been used in the literature to define energy functions [18], to characterize a signed graph's stability. Balanced triads (where the product of signs is positive) are considered stable and having negative energy, while unbalanced triads are considered unstable. Here, we consider the un-normalized [3] energy of a set of nodes C as the number of balanced triads minus the number of unbalanced triads.

$$\nu(C) = \sum_{\{v_i, v_j, v_k\} \in T(C)} -s_{ij} s_{jk} s_{ik} \tag{1}$$

where $T(C)$ is the set of triads in C, and s_{ij} represents the sign of the edge between i and j. We can easily see that the marginal contribution of a node i to C would be only dependent on the number of triangles v_i is a part of. If the other two nodes of a triad to which i belongs are already present in C, then i can make a marginal contribution through the triad. For a given pair of adjacent neighbors j and k of a node i, this will only happen in $\frac{1}{3}$ of the permutations. Thus, the shapley value $SV(v_i)$ will be given by

$$SV(v_i) = \sum_{\{v_j, v_k\}, v_j \in N(v_k), v_j, v_k \in N(v_i)} \frac{1}{3}(-s_{ij} s_{jk} s_{ik}) \tag{2}$$

The time complexity T_{SBSV} will be $O(\sum_{v_i \in V} (d^{total}_{in}(v_i))^2)$.

[3] We do not normalize the energy by the largest possible number of triads $\binom{|C|}{3}$, or the total number of triangles, since this makes computing the closed form difficult

Table 1: Comparison of Measures For Troll Detection

Approach	In Top 96	AP	MAP-5	MAP-10	MAP-20
FMF	10	0.031	0.031	0.033	0.033
NTV	7	0.044	0.043	0.045	0.042
NPF	18	**0.104**	0.104	0.107	0.108
FAT	15	**0.092**	0.092	0.094	0.093
NFADT	**19**	**0.130**	0.131	0.134	0.133
2 Hop FMF	**25**	0.158	0.157	0.153	0.155
3 Hop FMF	7	0.023	0.025	0.025	0.025
2 Hop NPF	16	0.148	0.139	0.143	0.152
3 Hop NPF	**22**	**0.193**	0.193	0.193	0.192
WBSV	4	0.039	0.039	0.039	0.038
SBSV	4	0.024	0.024	0.024	0.023

3.4 Weak Balance Shapley Value (WBSV)

Weak balance is a variant of balance theory [16], according to which the triad with three negative signs (all nodes are mutual enemies), is not considered as unbalanced. As a result, triads where s_{ij}, s_{jk} and s_{ik} are all -1 are not included in the value function, and consequently, in the Shapley Value. The time complexity will be the same as that of WBSV.

3.5 Evaluation and Results

We consider the task of ranking users to detect trolls in the Slashdot signed network [4], with 96 users annotated as trolls. The network has 71512 nodes, with 487751 edges. About 24% of these edges are negative.

We evaluate a centrality measure by first ranking the nodes of the graph in ascending order according to them, and then evaluating these ranklists based on how high the ground truth trolls rank in them. For evaluation, we consider two metrics

1. The number of trolls in the top g elements of the ranklist, where g is the number of ground truth trolls
2. The average precision (AP) metric from IR [19], with the trolls corresponding to "relevant documents". The metric is defined as follows

$$AP = \sum_{k=1}^{k=n} \frac{P(k) \times Mal(k)}{g} \tag{3}$$

Here, $P(k)$ is the precision (fraction of ground truth trolls) in the top k elements of the ranklist. $Mal(k)$ is 1 if the kth user is a troll, while g is the number of ground truth trolls. n is the size of the ranklist.

Besides computing these for the full graph, we also compute the mean of the AP (MAP) over 50 subgraphs formed by deleting 5 %, 10 % and 20% of the nodes. We observe that the **NPF**, **FAT** and **NFADT** measures perform considerably better than **FMF**, both on the full graph as well as on each of the random subsamples. Moreover, the **3-Hop NPF** gives the highest average precision amongst all the measures, while the performance of **k-Hop FMF** decreases when we go to 3 hops.

4 Conclusion

We propose here, for the first time, cooperative game theoretic Shapley Value based centrality measures for a signed social network. Furthermore, we demonstrate that Shapley Value based generalizations of even a simple measure such as FMF can detect trolls more effectively.

[4]http://konect.uni-koblenz.de/networks/slashdot-zoo

5 Future Work

The set of measures we propose here are based on notions of degree or number of k-hop neighbors, for a subset of nodes. In the future, we could look at exploring other notions, such as community score (e.g modularity) of a subset. Moreover, we could look at evaluating the proposed measures for other centrality-based ranking tasks, such as retrieving relevant publications for understanding a given publication, based on a signed citation network (Citations could be thought of as positive or negative based on whether the cited publication is criticized or praised).

References

[1] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. Signed networks in social media. In *SIGCHI*, 2010.

[2] Jérôme Kunegis, Andreas Lommatzsch, and Christian Bauckhage. The slashdot zoo: mining a social network with negative edges. In *WWW*, 2009.

[3] Paolo Massa and Paolo Avesani. Controversial users demand local trust metrics: An experimental study on epinions. com community. In *Proceedings of the National Conference on artificial Intelligence*, volume 20, page 121. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2005.

[4] Silviu Maniu, Bogdan Cautis, and Talel Abdessalem. Building a signed network from interactions in Wikipedia. In *Databases and Social Networks*, 2011.

[5] Efthymios Kouloumpis, Theresa Wilson, and Johanna Moore. Twitter sentiment analysis: The good the bad and the omg! *ICWSM*, 11:538–541, 2011.

[6] Kazuya Okamoto, Wei Chen, and Xiang-Yang Li. Ranking of closeness centrality for large-scale social networks. In *Frontiers in Algorithmics*, pages 186–195. Springer, 2008.

[7] Ulrik Brandes. A faster algorithm for betweenness centrality*. *Journal of Mathematical Sociology*, 25(2):163–177, 2001.

[8] Mark EJ Newman. A measure of betweenness centrality based on random walks. *Social networks*, 27(1):39–54, 2005.

[9] Phillip Bonacich. Some unique properties of eigenvector centrality. *Social Networks*, 29(4):555–564, 2007.

[10] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: bringing order to the web. 1999.

[11] Srijan Kumar, Francesca Spezzano, and VS Subrahmanian. Accurately detecting trolls in Slashdot Zoo via decluttering. In *ASONAM, 2014*, 2014.

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

[13] N Rama Suri and Yadati Narahari. Determining the top-k nodes in social networks using the Shapley value. In *AAMAS*, 2008.

[14] Karthik V Aadithya, Balaraman Ravindran, Tomasz P Michalak, and Nicholas R Jennings. Efficient computation of the shapley value for centrality in networks. In *Internet and Network Economics*. 2010.

[15] Ramasuri Narayanam, Oskar Skibski, Hemank Lamba, and TP Michalak. A Shapley value-based approach to determine gatekeepers in social networks with applications. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI14)*, 2014.

[16] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. Predicting positive and negative links in online social networks. In *Proceedings of the 19th international conference on World wide web*, pages 641–650. ACM, 2010.

[17] Dorwin Cartwright and Frank Harary. Structural balance: a generalization of heider's theory. *Psychological review*, 63(5):277, 1956.

[18] Seth A Marvel, Steven H Strogatz, and Jon M Kleinberg. Energy landscape of social balance. *Physical review letters*, 103(19):198701, 2009.

[19] Marc A Najork, Hugo Zaragoza, and Michael J Taylor. HITS on the web: How does it Compare? In *SIGIR*, 2007.