Unveiling Polarization in Social Networks: A Matrix Factorization Approach

Md Tanvir Al Amin*, Charu Aggarwal†, Shuochao Yao*, Tarek Abdelzaher*, and Lance Kaplan‡

*University of Illinois at Urbana-Champaign, Urbana, IL 61801, {maamin2, syao9, zaher}@illinois.edu
†IBM T. J. Watson Research, Yorktown Heights, NY 10598, charu@us.ibm.com
‡Networked Sensing and Fusion Branch, U.S. Army Research Laboratory, Adelphi, MD 20783, lance.m.kaplan.civ@mail.mil

Abstract—This paper presents *unsupervised* algorithms to uncover polarization in social networks (namely, Twitter) and identify polarized groups. The approach is language-agnostic and thus broadly applicable to global and multilingual media. In cases of conflict, dispute, or situations involving multiple parties with contrasting interests, opinions get divided into different camps. Previous manual inspection of tweets has shown that such situations produce distinguishable signatures on Twitter, as people take sides leading to clusters that preferentially propagate information confirming their individual cluster-specific bias. We propose a model for polarized social networks, and show that approaches based on factorizing the matrix of sources and their claims can automate the discovery of polarized clusters with no need for prior training or natural language processing. In turn, identifying such clusters offers insights into prevalent social conflicts and helps automate the generation of less biased descriptions of ongoing events. We evaluate our factorization algorithms and their results on multiple Twitter datasets involving polarization of opinions, demonstrating the efficacy of our approach. Experiments show that our method is almost always correct in identifying the polarized information from real-world twitter traces, and outperforms the baseline mechanisms by a large margin.

I. INTRODUCTION

This paper presents algorithms to uncover polarization on social media networks, such as Twitter, and identify opposing sets of biased tweets. We define polarization as a condition in which two opposing views enjoy wide support by different groups in a community. For example, a community might become divided over a political or social issue; this is often manifested as opposing views on how the issue should be resolved. Often, the conflict extends to claims about factual observations, such as whether a person had a gun on them or not at a particular time. These widely held and reported conflicting beliefs obfuscate descriptions of the real progression of events as a result of various injected biases. To uncover a less biased (i.e., more neutral) description of events, it is important to identify polarization and distill neutral observations from the reported mix, which motivates the work in this paper.

In this paper, we present a polarization model for information networks, and show that the presence of polarized groups can be detected by considering dependence among posted observations. Using matrix rank as a parameter, we propose a matrix factorization approach to uncover polarization. We explore different degrees of polarization and compare the quality of separation (of tweets of opposing polarity) across different algorithms using real traces collected from Twitter. The work is motivated, in part, by the increased reliance on social networks as news sources. Social network based news dissemination is different from traditional news networks, where raw information goes through curation by expert analysts and journalists before publication. In contrast, in the social media, anybody can post anything. Polarization or bias is inevitable [1]. Hence, tools are needed to clean-up the media before consumption as news.

Our results demonstrate that opposing sets of polarized tweets and sources can be identified automatically (with no content analysis or natural language processing) by the aforementioned matrix factorization approach. Experiments show that the proposed algorithm performs much better than the baseline methods in terms of accuracy in unveiling the polarized sources and groups. The underlying intuition lies in that, in cases of conflict, parties of opposing views tend to disseminate dissimilar sets of claims on social networks. Hence, some of the disseminated tweets can be separated into two subsets propagated by largely non-overlapping sets of sources. We represent the set of tweets as a matrix, where one dimension represents sources and the other represents their tweets (or claims), and where non-zero entries represent who said/forwarded which tweet. Given such a matrix, our algorithm uncovers the underlying latent groups and claims, thereby identifying both the conflicted social groups and their respective views. The language-independent nature of the approach makes it especially advantageous in applications involving multilingual media such as Twitter, since no dependency on a particular lexicon is involved.

It should be noted that the problem addressed in this paper is different from detecting communities in the social network. We observe that in practice, the crawled traces of polarized tweets are often intermixed with a large set of neutral sources and observations. Since neutral observations can also be relayed by polarized sources, and since neutral sources may mix and match view of different polarities, the task of separating the polarized clusters becomes a much harder problem. In this setting, algorithms based on community detection do not correctly identify the polarized clusters. In this paper, solutions are presented that explicitly handle the existence of a neutral population of sources and claims that blurs the boundaries of polarized groups.

The problem addressed in this paper is also different from the commonly addressed problem of veracity analysis on

social media. There can be several types of bias present on the social medium. In the extreme case, one or more of the polarized sources are malicious. People post false information to glorify or defame certain acts or causes. This propaganda may result in an 'online war' on the social platform, and veracity analysis might be used to detect improbable claims. However, it is often that polarization is more benign. Individuals do not post entirely fabricated observations, but rather color true observations depending on their opinions. A more subtle form of polarization occurs when people selectively propagate or suppress observations based on their bias. For example, a person supporting political party X may only forward (true) positive information about X and the negatives about Y. Another person can forward (true) information about the opposite. In this case, veracity analysis does not help identify polarization.

Also, note that the problem is different from sentiment analysis. A statement that mentions, say, a president and features a negative sentiment might not actually be opposing the president. It might be negative on something else. For example, consider these tweets regarding a former Egyptian president (Morsi): "Saudi Arabia accused of giving Egypt $1B to oust Morsi" or "Egypt clashes after army fire kills #Morsi supporters". Both tweets mention "Morsi" (the president) and feature negative sentiments (due to use of such keyword as "accuse", "oust", "clash" and "kill"). However, reading them carefully, it is easy to see that both sympathize with the president depicting him and his supporters as victims.

The paper is therefore novel in addressing the problem of identifying and separating *polarization* as opposed to, for example, performing veracity analysis, community detection, or sentiment analysis.

The rest of this paper is organized as follows. Section II illustrates a motivating example from Twitter that leads to our problem, models polarization in social network, and formulates the problem. Section III derives a matrix factorization based gradient descent algorithm to estimate the polarities. In section IV we describe the implementation, evaluate our algorithm, and compare it to other baselines. Section V reviews the literature related to polarization in social networks. The paper concludes with a discussion in section VI.

II. POLARIZATION IN SOCIAL NETWORKS

The end result (of identifying polarization), presented in this paper, could in principle be accomplished using semantic analysis and natural language processing. The goal of our work, however, is to achieve that end in a *language agnostic* manner. There are two reasons why this is important. First, on a multi-lingual, multi-national medium, such as Twitter, the number of languages used is large. Developing a model for each language specifically to identify polarization is a rather expensive undertaking. Second, it is not always clear that understanding the language helps understand the polarity of a statement. Consider, for example, the following tweet about Jamala, the winner of the Eurovision competition in 2016: "Jamala performs Bizim Qirim at Kiev concert Hall,

18 May, 2015. The same song wins Eurovision one year later". Is this tweet advertising Jamala (i.e., is "pro") or is it against her (i.e., is "anti")? Someone not familiar with the underlying background might consider it pro. In reality, it is not. Eurovision rules dictate that Eurovision songs have to be original. By claiming that the song was performed a year earlier, the source suggests that the entry should have been disqualified. The need to understand situation-specific context on a case-by-case basis poses significant challenges when it comes to building general-purpose schemes for identifying polarity.

Our approach uses a different intuition. Individuals retweeting statements such as the above, on average, understand their context and polarity. Their behavior reflects their understanding. Hence, by monitoring such collective behavior (namely, the overall propagation patterns of tweets), and clustering it by similarity, it is possible to separate "pro" versus "anti" without having to understand the language. In essence, we harness the collective intelligence of the social medium. In the following sections, we introduce the information model for polarized social networks, and formally define the problem.

A. Information Model for Social Networks

Online social platforms often allow mechanisms to crawl public information. The crawled information at first goes through domain specific cleaning or filtering steps. The content is then clustered using appropriate similarity measurement, which helps to consolidate small variations in the data, and generate a rich information network. A cluster of the very similar observations is considered as a single *assertion*, and the people or the authors who posted those observations are considered as *sources*. The bipartite graph from the *sources* to the *assertions* is called a *source-assertion network*. The method of generating this network from the crawled data has been discussed in detail in different works [2], [3]. In this paper, we represent the source-assertion network as a binary *source-assertion matrix* A of dimensions $s \times c$, where s is the number of sources, and c is the number of assertions. If source i claims assertion j, then $a_{ij} = 1$, otherwise $a_{ij} = 0$.

In addition to the source-assertion network, a social influence or dependency network can also be derived (or crawled), where an (s, t) edge denotes that source s has a tendency to forward information if it is received from source t. This graph can be weighted when the intensity of influence or dependency is considered into the model, or it can be simplified as an unweighted graph of binary relations. We represent it as a $s \times s$ *social dependency matrix* $T = [t_{ij}]$ of binary values. It can be derived from an explicit social network such as Twitter follower-followee relations. It can also be estimated from retweet behavior of the sources. Netrapalli and Sanghavi [4] model the propagation of information through the social network as cascades of epidemics. Given the tweets along with sources and timestamp information, they solve the inverse problem of finding the latent propagation structure. In this paper, we estimate the social dependency network using their maximum likelihood estimation mechanism.

B. Modeling Polarized Information Networks

In this section we augment the source-assertion network with additional states for polarized scenario. Please note that these models are developed according to the real-world observations reported by multiple independent works [5], [6].

We define a *polarity group* as the set of different senses (*polarities*) relative to a polarity context (*pole*). For example, in a US political parties context, the polarity group can be $K = \{$democrat, republic$\}$. Please note that although two polarities are common, the polarity group can contain more than two members, if the context is not bipartite. For example, the polarity group of the former example could also be $K' = \{$democrat, libertarian, republic$\}$. Given a set of assertions strictly related to the polarity context, each assertion can be classified as one of the polarities from the polarity group. However, due to the nature of data collection, often there are assertions that do not belong to any of the polarities, which can be termed as neutral, or *nonpolarized assertions*. For example, every tweet that contains the keyword Morsi is not necessary pro-Morsi or anti-Morsi.

Obtaining the ground truth about the polarity of an assertion requires human effort (that we want to automate). It requires a human grader to understand the content of the assertion and its context. Then the grader assigns a polarity from the polarity group, or classifies it as a neutral or nonpolarized assertion. A source is polarized if its odds of making non-neutral claims of a particular polarity is above a threshold τ, otherwise a source is neutral.

Suppose the polarity group is $K = \{$pro, anti$\}$. The bipartite source-assertion network takes the form shown in Figure 1a. Circles $S1$ to $S6$ represent six sources, and squares $C1$ to $C7$ represent seven assertions. Empty circles (squares) represent neutral sources (assertions). Filled circles (squares) represent polarized sources (assertions). Arrows represent claims of different polarities. The relationship between the polarized and the neutral components can be represented as Figure 1b. Here sources (assertions) with particular polarities are consolidated together as a single circle (square) representing that polarity. The rest of the vertices are consolidated as neutral sources and neutral assertions. The polarized vertices are consolidated as the polarized network.

C. Problem Formulation

Consider a scenario with two polarities, namely pro and anti. Some of the observations posted in the social media favors or averts the pro camp. Some of the observations do the same for the anti camp. In this paper, our goal is to separate the polarities. To develop the formulation, we consider the simplified case with opposing polarities only, without the presence of the neutral network. Later we show how the solution to the simplified formulation is adapted to solve the general case with a huge neutral network obscuring the polarized network.

The observation that there are polarized factions that do not share posts contradicting their polarities, allows us to separate them. Consider $\Pr(S_iC_j)$, probability that source i shares an
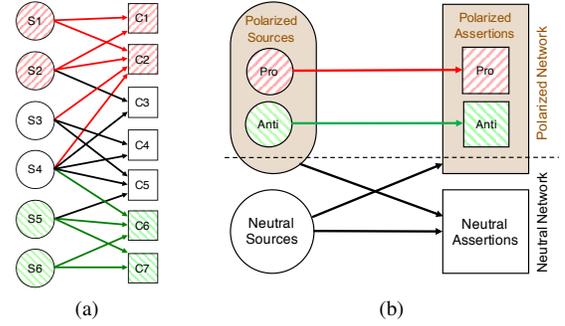

Fig. 1. (a) Model of a polarized source-assertion network, (b) Relation between the polarized and the neutral network.

observation j. $\Pr(C_j^q)$ denotes the probability that assertion j is of polarity q. $\Pr(S_i^q)$ denotes the probability that source i is of polarity q. We can then write equation 1.

$$
\begin{aligned}
\Pr&(S_iC_j) \\
&= \Pr(S_iC_j|C_j^{pro}).\Pr(C_j^{pro}) + \Pr(S_iC_j|C_j^{anti}).\Pr(C_j^{anti}) \\
&= \Pr(S_iC_j|S_i^{pro}C_j^{pro}).\Pr(S_i^{pro}).\Pr(C_j^{pro}) \\
&\quad + \Pr(S_iC_j|S_i^{anti}C_j^{pro}).\Pr(S_i^{anti}).\Pr(C_j^{pro}) \\
&\quad + \Pr(S_iC_j|S_i^{pro}C_j^{anti}).\Pr(S_i^{pro}).\Pr(C_j^{anti}) \\
&\quad + \Pr(S_iC_j|S_i^{anti}C_j^{anti}).\Pr(S_i^{anti}).\Pr(C_j^{anti}) \quad (1)
\end{aligned}
$$

When the assertion opposes the polarity of a source, the source is not going to share it. Therefore, both $\Pr(S_iC_j|S_i^{pro}C_j^{anti})$ and $\Pr(S_iC_j|S_i^{anti}C_j^{pro})$ reduces to 0.

$$
\begin{aligned}
\Pr(S_iC_j) = \;&\Pr(S_iC_j|S_i^{pro}C_j^{pro}).\Pr(S_i^{pro}).\Pr(C_j^{pro}) \\
&+ \Pr(S_iC_j|S_i^{anti}C_j^{anti}).\Pr(S_i^{anti}).\Pr(C_j^{anti}) \quad (2)
\end{aligned}
$$

Now we consider the terms $\Pr(S_iC_j|S_i^{pro}C_j^{pro})$ and $\Pr(S_iC_j|S_i^{anti}C_j^{anti})$ in equation 2. In an ideal situation, these values are 1, making a source share each and every observation whenever the polarity matches. In practice, this does not happen. $\Pr(S_iC_j|S_i^qC_j^q)$ depends on various social and human factors, but we can simplify this probability as a combination of two *independent* components, (i) activity level of the source i denoted by $act(S_i)$, and (ii) circulation level of the assertion j denoted by $cir(C_j)$. Taking δ as a scaling constant, we can write $\Pr(S_iC_j|S_i^qC_j^q) = \delta.act(S_i).cir(C_j)$.

$$
\begin{aligned}
\Pr(S_iC_j) = \;&\delta.act(S_i).cir(C_j).\Pr(S_i^{pro}).\Pr(C_j^{pro}) \\
&+ \delta.act(S_i).cir(C_j).\Pr(S_i^{anti}).\Pr(C_j^{anti}) \,(3)
\end{aligned}
$$

Consider the general case with k polarities, $q \in \{1..k\}$. $U = [u_{iq}]$ is an $s \times k$ matrix, and $V = [v_{jq}]$ is a $c \times k$ matrix. Activity levels of the sources and their probabilities to belong to particular polarized camps are represented in U. Circulation levels of the assertions and their probabilities to favor particular camps are represented in V. Therefore, $u_{iq} = \delta_1.act(S_i).\Pr(S_i^q)$, and $v_{jq} = \delta_2.cir(C_j).\Pr(C_j^q)$. If $\hat{A} = [\hat{a}_{ij}]$ represents the probability of a source to share a particular assertion, we can rewrite equation 3 as $\hat{a}_{ij} = \sum_{q=1}^k u_{iq}v_{jq}$, or $\hat{A} = UV^T$.

Given $A = [a_{ij}]$ as the actual observations on whether source i shared assertion j in the social network, $T = [t_{ij}]$ as the social dependency matrix on whether source i is likely to forward information received from source j, and a polarity group K, our goal is to estimate U and V component matrices that allow us to separate the polarized components.

D. Solution Approach

Given k as the rank, we can factorize A to estimate U and V components. Please note that, $A = UV^T = URR^{-1}V^T = (UR)(VR^{-T})^T$, where R is a $k \times k$ multiplier matrix. Therefore, factorizing A without any constraint will result in UR and VR^{-T} as component matrices. In the following section, we add appropriate constraints to limit the arbitrariness of R. In the simplified case, when there is only the polarized network, rank of A is exact. Sources and assertions of different polarities can be uniquely separated using the estimated factor matrices $\hat{U}$ and $\hat{V}$. However, in the presence of a large neutral network, the number of polarized camps k does not correctly represent the rank of A. In this case, different separations are possible that can approximate the observation matrix A. We estimate multiple instances of $(\hat{U}, \hat{V})$ using different initializations. For each instance, observations are partitioned into different polarities. Instances are generally related to each other in terms of similarity between corresponding partitions. Anomalous instances that are highly different than the rest are discarded. Rest of the instances are aggregated to estimate the final partitions.

III. A MATRIX FACTORIZATION APPROACH TO UNCOVER POLARIZATION

In this section, we derive a gradient-descent algorithm to jointly estimate the polarization of the sources and assertions. Suppose A is the $s \times c$ source-assertion matrix. Polarization of the sources and the assertions can be estimated from A by factorizing it in the form of matrices $\hat{U}$ and $\hat{V}$, defined earier.

If $k = rank(A)$, A can be factorized exactly in the form $A = UV^T$, where $U = [u_{ij}]$ is an $s \times k$ matrix that represent the polarization of the sources, and $V = [v_{ij}]$ is a $c \times k$ matrix representing the polarization of the assertions. Please note that A is an incomplete matrix because when source i does not claim assertion j, it can be that source i did not have opportunity to observe j, or i ignored assertion j after observing it. Therefore a sample of the missing edges in the source-assertion network are represented as $a_{ij} = 0$, and the rest are considered as missing. Because A is incomplete, we do not know the exact rank of A. However, as visible from Figure 1a, the sources and assertions of different polarized camps are independent when they are sharing information related to the particular polarized scenario. Hence, we take the number of polarized groups $|K|$ as $rank(A)$, and *approximately* factorize A as $A \approx UV^T$.

Note that this condition is defined only for the entries of A that are observed. Therefore, let us define the set $O = \{(i,j) : a_{ij} \text{ is observed}\}$ to be all the indices in matrix A that are observed. Given a particular U and V, the estimate of an entry of A is given by $\hat{a}_{ij} = \sum_{q=1}^k u_{iq}v_{jq}$. Therefore, the estimation error is $e_{ij} = a_{ij} - \hat{a}_{ij}$. In order to approximately factorize the matrix A, we need to minimize the objective function, which is equal to the sum-of-squared errors $J = \sum_{(i,j) \in O} e_{ij}^2 = \sum_{(i,j) \in O} (a_{ij} - \sum_{q=1}^k u_{iq}v_{jq})^2$. This form of the objective function, however, can result in infinitely many solutions, each of which minimizes J. Therefore, we impose the following constraints. The first constraint corresponds to overfitting of the objective function. The second constraint corresponds to the impact of the social dependency matrix.

1) Regularization: If U and V is a particular solution, then multiplying U by an arbitrary $k \times k$ real matrix R, and multiplying V by R^{-T} would also minimize J, provided R is inversible. This is because $UV^T = UIV^T = URR^{-1}V^T$, where I is a $k \times k$ identity matrix. Depending on the chosen initial values or the missing entries, the objective function can overfit the model, or oscillate between multiple solutions. Therefore, we impose a regularization constraint on J. We choose to use L_2-regularization $\lambda(||U||_F^2 + ||V||_F^2)$ so that arbitrarily large values in R would be prevented. The multiplier $\lambda > 0$ represents the value of the regularization parameter.

2) Social dependency-based polarization consistency: We observed that polarization in the crawled data is obscured by a large nonpolarized or neutral network. Presence of such sources and assertions result in multiple separations between the different polarity groups likely. The objective function would result in multiple candidate solutions. Therefore, we add an additional constraint. Users that depend on one another according to the social dependency matrix T are more likely to exhibit polarization consistency. So the columns in U corresponding to sources who depend on each other contains similar entries. If $\overline{u_x}$ is the row in U corresponding to source x, the additive component $\gamma t_{ij}||\overline{u_i} - \overline{u_j}||^2$ would add a penalty whenever source i depends on source j, but their corresponding columns vary. Here $\gamma > 0$ is a parameter that regulates the importance of the social consistency component. This parameter can be chosen later in the tuning phase. Please note that adding this constraint will increase the error in the factorization but it will favor solutions that have higher consistency with the social dependency network.

Therefore, by adding these terms, our objective function becomes $J = \sum_{(i,j) \in O} (a_{ij} - \sum_{q=1}^k u_{iq}v_{jq})^2 + \sum_{i,j} \gamma t_{ij}||\overline{u_i} - \overline{u_j}||^2 + \lambda(||U||_F^2 + ||V||_F^2)$, which needs to be minimized.

A. Solving the Optimization Problem

We minimize J with respect to the parameters in U and V using gradient descent method. We rewrite the objective function, and compute the partial derivative of J with respect to each parameter in U and V:

$$J = \sum_{(i,j) \in O} e_{ij}^2 + \sum_{i,m} \gamma t_{im} \sum_{q=1}^k (u_{iq} - u_{mq})^2$$
$$+ \lambda(||U||_F^2 + ||V||_F^2) \qquad (4)$$
$$\frac{\partial J}{\partial u_{iq}} = 2 \sum_{j:(i,j) \in O} e_{ij}(-v_{jq})$$

```
1: procedure FACTORIZE(A, T, k)
2:     Randomly initialize U, V
3:     repeat
4:         for each (i, q) do
5:             u⁺_{iq} ← u_{iq} − α ∂J/∂u_{iq}          ▷ Equation 5
6:         end for
7:         for each (j, q) do
8:             v⁺_{jq} ← v_{jq} − α ∂J/∂v_{jq}          ▷ Equation 6
9:         end for
10:        for each (i, q) do
11:            u_{iq} ← u⁺_{iq}
12:        end for
13:        for each (j, q) do
14:            v_{jq} ← v⁺_{jq}
15:        end for
16:    until convergence reached on U, V
17:    return (U, V)
18: end procedure
```

Fig. 2. Gradient descent algorithm for factorization

$$+ 2\sum_m (\gamma t_{im} + \gamma t_{mi})(u_{iq} - u_{mq}) + 2\lambda u_{iq} \quad (5)$$

$$\frac{\partial J}{\partial v_{jq}} = 2 \sum_{i:(i,j)\in O} e_{ij}(-u_{iq}) + 2\lambda v_{jq} \quad (6)$$

Note that we can ignore the constant factor of 2 throughout the RHS of the aforementioned equation for the purposes of gradient descent. We compute all partial derivatives with respect to the different parameters in u_{iq} and v_{jq} to create gradient matrix ∇U of dimensions $s \times k$, and ∇V of dimensions $c \times k$. The gradient-descent method updates $U \Leftarrow U - \alpha\nabla U$, and $V \Leftarrow V - \alpha\nabla V$, where α is the step-size. The parameter γ, λ can be selected using cross-validation. Figure 2 enumerates this mechanism.

We impose the additional constraint that the entries of the matrix U and V are non-negative, although the optimization objective function remains the same. It provides a sum-of-parts decomposition to the source-assertion matrix as dictated by the problem formulation. To achieve this, during initialization, the entries of matrices U and V are set to non-negative values in $(0, 1)$. During an update, if any entry in U or V becomes negative, then it is set to 0.

B. Separating Polarities using $\hat{U}$ and $\hat{V}$

Activity levels of the sources and their probabilities to belong to particular polarized camps are represented in U. Circulation levels of the assertions and their probabilities to favor particular camps are represented in V. Rows of U and V can be considered as points in a k-dimensional euclidean space. In the simplified case, where the source-assertion matrix consists of only the polarized network with $K = \{\texttt{pro}, \texttt{anti}\}$, the extreme points of U or V are $(1, 0)$ or $(0, 1)$. These points represent the sources making all the $\texttt{pro}$ assertions, or making all the $\texttt{anti}$ assertions, respectively. All the other points would fall on either x-axis or y-axis. However, in the general case, the neutral network is present, hence it is possible to have points that fall within the right triangle defined by vertices at $(0, 0)$, $(1, 0)$, and $(0, 1)$.
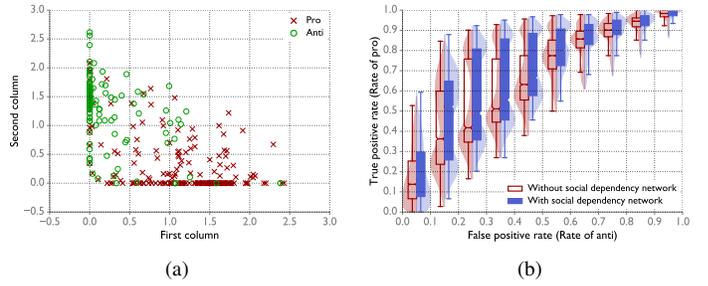

Fig. 3. (a) Assertions from the estimated factor matrix $\hat{V}$ and their polarities, (b) Although the social dependency network improves performance, there is still variance in the separation due to the presence of the neutral network.

Through factorization we have estimated $\hat{U} = UR$, and $\hat{V} = VR^{-T}$. This multiplier R causes the estimated values in $\hat{U}$ or $\hat{V}$ to have been applied a linear transformation. A linear transformation in general can be decomposed to several rotations and scales. Due to the constraints we have added to J, effect of R is small. Figure 3a shows an output where the rows of $\hat{V}$ are plotted on the $2D$ plane for a particular experiment. We observe that the multiplier R has been mostly restricted to a diagonal matrix corresponding to scale transformation.

Figure 3a also plots the ground truth of the assertions as obtained via manual annotaion. To separate the different polarity groups, we note that linear transformations preserve parallel lines. Therefore, the midpoint of a transformed line corresponds to the transformation of the midpoint of the original line. We can separate the polarities by finding the pair of assertions (a, b) from $\hat{V}$ with maximum euclidean distance, i.e. $\arg\max_{a,b} ||\hat{v}_a - \hat{v}_b||^2$, and assigning the other assertions to either the polarity of a or b, using a nearest neighbor rule. However, we observe that R has been mostly restricted to scaling. Therefore, to obtain a separation of the polarities, assertion j can be assigned to the group corresponding to $\arg\max_q \{\hat{v}_{jq}\}$. For Figure 3a, this corresponds to using sign of $\hat{v}_{j,1} - \hat{v}_{j,2}$ as the separator. The sources can also be separated in a similar manner using $\hat{U}$.

C. Ensemble of Factorization Experiments

We note that in the presence of a large neutral network, different runs of factorization results in different separations. Figure 3b illustrates the receiver operating characteristics (ROC) for the egypt scenario. ROC curve plots true positive rate vs. false positive rate, and is used to assess the quality of classification. The optimal algorithm has an area of 1 under the ROC, which happens when the output includes all the true positives before any of the false positives.

Figure 3b plots the distribution of true positive rate for different false positive rates, and shows that although the factorization algorithm is able to achieve good performance, there is significant variance in the separation obtained from the results. We also compare the result of when the social dependency network is used as a constraint vs. when it is not. We observe that although use of social dependency network

improves the quality of the results, there is still variance in the separation. We, therefore, use an ensemble of factorization experiments to estimate the most likely assignments of the assertions to the respective polarities.

It is not possible to directly compare $\hat{V}_m$ with $\hat{V}_n$, when m and n different experiments, because of the transformation difference caused by R_m and R_n. We, therefore, separate the assertions to different polarity groups for each experiment. Experiments are aligned to each other using a mechanism based on Jaccard distance [7]. We explain it for two polarities, i.e. $k = 2$. Figure 4 shows the algorithm, with the procedure ESTIMATEPOLARITIES at line 38 being the starting point. Suppose the separation generated by factorization experiment m is B_m^1 and B_m^2, and the separation generated by factorization experiment n is B_n^1 and B_n^2. It is possible that (B_m^1, B_m^2) aligns with (B_n^1, B_n^2), or with (B_n^2, B_n^1). Figure 5a illustrates the two cases considering the polarities as pro and anti. We compute a 2×2 matrix of the Jaccard distances between the separations created by the experiments. Jaccard distance between two sets X, Y is defined as $1 - \frac{|X \cap Y|}{|X \cup Y|}$. It is used to assess how similar or dissimilar they are. In order for the two experiments to match, either the main diagonal will exhibit more similarity than the anti-diagonal (or vice versa). If the maximum in the matching diagonal is below a threshold τ_{edge}, given their difference is within τ_{diag}, the experiments are considered to match and a weighted edge is added to G, the graph of experiments. The weight is considered positive if the experiments matched along the main diagonal, and negative if the experiments matched along the anti-diagonal. Figure 5b shows an experiment graph (without the weights) obtained from 20 experiments on the egypt polarized scenario. Experiments that highly differ from the others remain isolated in the experiment graph, or form small islands. We find the experiment with the largest degree in G, and agreegate all the adjacent experiments.

There can be several procedures to aggregate the experiments. We keep a vector of frequencies (x_j, y_j) for each assertion. x_j and y_j counts the number of times assertion j has been assigned to polarity x or y. Normalizing these frequencies and sorting them by the difference of the vector components $(x_j - y_j)$ gives us a spectrum of assertions, from the most likely to belong to one polarized camp to the most likely to belong to the other camp.

IV. EVALUATION

We evaluate our algorithm in the context of polarized scenario in Twitter. Tweets were crawled in real time with tools using Twitter search API. Three sets of traces were collected that contains polarization around (i) former Egyptian president Morsi, (ii) Eurovision song contest 2016 winner Jamala, (iii) US Presidential election candidate Donald Trump. The entire collection of recorded traces was clustered based on text similarity to generate a representative summary [3], [7], [8]. We implemented the factorization program using Java. Sparse matrix data structures were used to efficiently store large matrices. Different components of the pipeline were

```
1:  procedure PARTITIONDISTANCE(B_m, B_n)
2:      ▷ JACCARDDIST(X, Y) = 1 − |X∩Y|/|X∪Y|
3:      d11 ← JACCARDDIST(B_m^1, B_n^1)
4:      d12 ← JACCARDDIST(B_m^1, B_n^2)
5:      d21 ← JACCARDDIST(B_m^2, B_n^1)
6:      d22 ← JACCARDDIST(B_m^2, B_n^2)
7:      dist ← 1.0
8:      if d11 < d12 and d22 < d21 and |d11 − d22| < τ_diag then
9:          ▷ (B_m^1, B_m^2) aligns with (B_n^1, B_n^2)
10:         dist ← max(d11, d22)
11:     else if d12 < d11 and d21 < d22 and |d12 − d21| < τ_diag then
12:         ▷ (B_m^1, B_m^2) aligns with (B_n^2, B_n^1)
13:         dist ← −max(d12, d21)
14:     end if
15:     return dist
16: end procedure

17: procedure GENERATEEXPGRAPH(B, size)
18:     G ← ∅
19:     for m ∈ [1, size − 1] do
20:         for n ∈ [m + 1, size] do
21:             dist ← PARTITIONDISTANCE(B_m, B_n)
22:             if |dist| < τ_edge then
23:                 ▷ Insert weighted undirected edge (m, n, dist) to G
24:                 G ← G ∪ (m, n, dist)
25:             end if
26:         end for
27:     end for
28:     return G
29: end procedure

30: procedure MERGEEXP(freq, B^X, B^Y)
31:     for v ∈ B^X do                    ▷ Assertions in B^X
32:         freq[v] ← freq[v] + (1, 0)
33:     end for
34:     for v ∈ B^Y do                    ▷ Assertions in B^Y
35:         freq[v] ← freq[v] + (0, 1)
36:     end for
37: end procedure

38: procedure ESTIMATEPOLARITIES(A, T, size)  ▷ Run size experiments
39:     ▷ A source-assertion matrix, T source dependency matrix, k = 2
40:     for l ∈ [1, size] do
41:         (Û_l, V̂_l) ← FACTORIZE(A, T, 2)           ▷ Figure 2
42:         (B_l^1, B_l^2) ← SEPARATEASSERTIONS(V̂_l)   ▷ Section III-B
43:     end for
44:     G ← GENERATEEXPGRAPH(B, size)     ▷ Graph of experiments
45:     node ← Vertex with maximum degree in G
46:     freq ← ∅        ▷ Mapping assertions to frequency of polarities
47:     MERGEEXP(freq, B_node^1, B_node^2)
48:     for each edge (node, exp, dist) ∈ G do
49:         if dist ≥ 0 then
50:             MERGEEXP(freq, B_exp^1, B_exp^2)
51:         else
52:             MERGEEXP(freq, B_exp^2, B_exp^1)
53:         end if
54:     end for
55:     prob ← ∅    ▷ Mapping assertions to distribution of polarities
56:     for each assertion v in freq do
57:         prob[v] ← (freq[v]_1/degree[node], freq[v]_2/degree[node])
58:     end for
59:     return prob
60: end procedure
```

Fig. 4. Algorithm to estimate polarities from an ensemble of experiments

interfaced using Python. Factorization was performed followed by the ensemble of multiple experiments to separate the tweets between two polarities. We used $k = 2$, $\alpha = 0.001$, $\gamma = 0.1$, $\lambda = 0.5$, ensemble $size = 20$, $\tau_{diag} = 0.15$, $\tau_{edge} = 0.7$. We
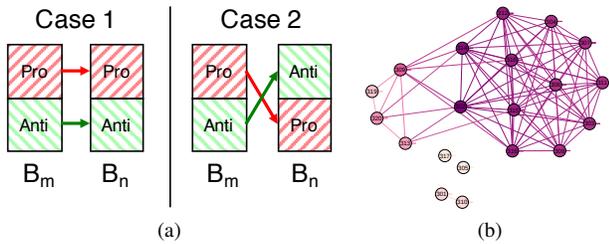
Fig. 5. (a) Aligning two experiments, (b) Graph of 20 experiments

TABLE I
TOP 5 TWEETS FROM THE SEPARATED POLARITIES (EGYPT)

	Pro-Morsi	Anti-Morsi
1	Sudden Improvements in Egypt Suggest a Campaign to Undermine Morsi http://t.co/0yCjbKGESr	Prayers for the Christian community in Egypt, facing violent backlash for opposing the Muslim Brotherhood. https://t.co/O5X7BwUjCI
2	Saudi Arabia accused of giving Egypt $1B to oust Morsi http://t.co/d4ZQNntCH	Egypt's Coptic Christians, under attack for supporting overthrow of Muslim Brotherhood, need continued prayers: http://t.co/dW0gdcielb
3	Before Morsi's Ouster, Egypt's top generals met regularly with opposition leaders http://t.co/LbdHKJF508 via @WSJ	Islamic extremists reportedly attacking Egypt's Christian community over Morsi ouster — Fox News http://t.co/VMMN2m49Sw
4	#Egypt: #Morsi supporters denied rights amid reports of arrests and beatings — Amnesty International http://t.co/koVRHlmdWk	In Egypt, the death toll in the clashes between police and pro-Morsi supporters in Cairo has risen to 34.
5	Crowds March in Egypt to Protest Morsi Detention http://t.co/Hp9566xyfB	Amnesty International — Egypt: Evidence points to torture carried out by Morsi supporters http://t.co/8hgAHrNoWd

compare the quality of separation obtained by our algorithm with the following related techniques:

1) Sentiment Analysis: Sentiment analysis [9], [10] uses language models to understand the sense of content written in natural language, and classifies them as having *positive*, *negative*, or *neutral* sentiment. To annotate the assertions we used Umigon [11] and Sentiment140 [12], two freely available specialized tools to perform sentiment analysis on tweets.

2) Community Detection: Polarized sources are unlikely to share tweets contradicting their own polarity. Therefore detecting communities in the social network is a candidate mechanism for separating the polarities. We partition the graph of sources and assertions into $k = 2$ communities with the objective of minimizing the *edge-cut* (number of edges that cross partitions). We used Metis [13] to obtain that. In addition to detecting communities, we have added another baseline where the assertions in each community are ranked by their degrees. We refer this mechasim by *MetisVoting*.

3) Veracity Analysis: Algorithms to perform veracity analysis [2], [14]–[16] utilize the source-assertion network to uncover likely facts from the set of tweets. They can be considered related techniques if one of the polarities have more affinity towards factual information. We used the EM-Social [2] algorithm to jointly asses the credibility of the sources and the assertions.

A. Egypt

Mass street protests against the then president Mohamed Morsi was followed by a coalition led by the army chief, on July 3, 2013. The president was deposed and arrested by the army along with other leaders of his political party. This incident resulted in protests and clashes between the supporters and the opponents of the removed president. Tweets related to the deposed president were collected. For the purpose of evaluation, the largest 1000 clusters containing English tweets were read and manually annotated on whether they were pro-Morsi, anti-Morsi, or neutral in sense. There were 199 pro-Morsi, 109 anti-Morsi, and 692 nonpolarized assertions.

Figure 6a compares the receiver operating characteristics (ROC) achieved from our algorithm to other baselines. To obtain the ROC, the set of assertions were sorted in the order of highest polarity in one class to the highest polarity in the other class. When consuming the assertions in that sequence, finding a pro-Morsi assertion was considered as an occurence of true positive, and finding an anti-Morsi assertion was considered

as an occurence of false positive. The area under ROC curve measures how well an algorithm performs both in terms of finding the correct answers, and omitting the wrong answers.

Factorization algorithm performs really well. Area under the ROC curve is approximately 0.93. Both Umigon and Sentiment140 performed just as good as a random technique, because (i) a large number of assertions were classified as neutral, and (ii) as described earlier in the paper, sentiment analysis is not the correct technique to uncover polarization. An assertion having positive sentiment can be a positive statement favoring either camp. Hence, sentiment is orthogonal to polarity. EM-Social is also unable to differentiate between the polarities. It illustrates that there was almost no correlation between the veracity of a tweet and any particular polarity. Metis and MetisVoting techniques performed better than the other baselines because of their graph partitioning nature. However, the source-assertion network had around 80% non-polarized sources and 70% nonpolarized assertions. Therefore a community detection analysis was unable to perform well.

Table I shows the top 5 tweets from each polarity from the separation achieved using our algorithm. Note that the tweets on the left column sympathize with the deposed president or his supporters. On the other hand, the tweets on the right column is vocal against the deposed president and his political party, and reporting negative news about them.

B. Eurovision

Susana Jamaladinova (Jamala) from Ukraine was the winner of Eurovision 2016, an annual European song competition. It was unexpected to many as the expected winner was Russia or Australia according to pre-competition polls. The winning song, 1944, according to the artist, was telling a personal story related to her family in the aftermath of the deportation of the Crimean Tatars by the Soviet Union. However, it was also alleged to have political connotations against Russian interference with Crimea in 2014. Tweets related to *Jamala* were collected for five days after her win. The largest 1000 assertions were manually annotated. There were 600 pro-Jamala, 239 anti-Jamala, and 161 neutral assertions.

Figure 6b compares quality of factorization with other baselines. Our algorithm performs best in this scenario. Metis performs reasonably better than the earlier case because of
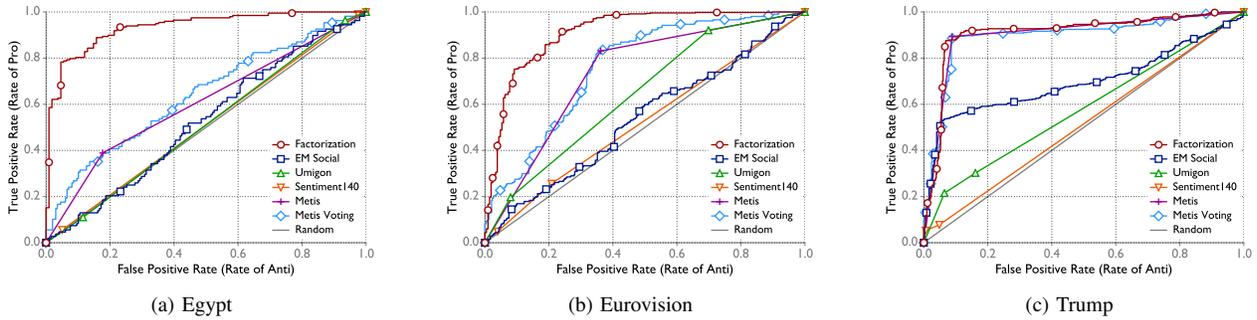
Fig. 6. (a) Egypt: Factorization performs best with area under ROC 0.93, EM Social 0.53, Umigon 0.51, Sentiment140 0.51, Metis 0.61, MetisVoting 0.64, (b) Eurovision: Factorization performs best with area under ROC 0.91, EM Social 0.54, Umigon 0.64, Sentiment140 0.52, Metis 0.73, MetisVoting 0.76, (c) Trump: Factorization 0.92, EM Social 0.70, Umigon 0.58, Sentiment140 0.52, Metis 0.90, MetisVoting 0.90

TABLE II
TOP 5 TWEETS FROM THE SEPARATED POLARITIES (EUROVISION)

	Pro-Jamala	Anti-Jamala
1	Incredible performance by #Jamala, giving Crimean Tatars, suffering persecution & abuse, reason to celebrate https://t.co/XWOZADrywH	For #Jamala1944: Crimea Tatar volunteers in the Nazi army parade before senior German officers, 1942 #Eurovision https://t.co/itgsyKvzO2
2	Breaking: #Russia launches harassment campaign against #Jamala's @Twitter a/c. All known Kremlin trolls. @BBC_ua @AP https://t.co/btk9QyUpkH	Repeat after me: NATO loves Jamala and there was absolutely nothing political about her win (via @marcelsardo) https://t.co/0bTbzVIR35
3	President awarded @jamala title of the Peoples Artist of Ukraine https://t.co/2df8J9zHP5	#BOOM Jamala released Eurovision song commercially on 19.06.2015 in Kiev club Atlas. @EBU_HQ https://t.co/LGbgb77RzH https://t.co/9Se1rwEkIg
4	Jamalas father: We do not talk with Russian journalists https://t.co/EEewVpZ9D2 https://t.co/51rcf8Je3K	#Oops Poroshenko accidently confirms on TV that Jamala's Eurovision song 1944 is the same song "Crimea is ours" from May 2015. @EBU_HQ
5	Congratulations to Ukraine on winning #Eurovision 2016! @JAMALA wrote and composed her song '1944' by herself. https://t.co/vZjYHvtoC	Second left grandfather Jamala!. Ordinary fascist, that "the tyrant Stalin sent him to Kyrgyzstan"! @antonio_bordin https://t.co/b9jsXHhiP1

TABLE III
TOP 5 TWEETS FROM THE SEPARATED POLARITIES (TRUMP)

	Pro-Trump	Anti-Trump
1	Retweet if you are 100 PERCENT voting for Donald Trump	Donald Trump said women should be punished for seeking an abortion. That's not a distraction – it's a disgrace. https://t.co/sbJ3opebyB
2	@realDonaldTrump Fugedaboudit!!! The woman in New York love Donald Trump!!! https://t.co/7yzgMHVzL4	At this point, Donald Trump has insulted the vast majority of Americans. The good news is, there's something we can all do about it: Vote.
3	1987: Donald J. Trump Celebrated As Model Citizen in #NYC. Remember TV without HD? #NYPrimary #MAGA #Trump2016 https://t.co/5ROvjhJyAK	Read and sign this letter that people all over are signing to Donald Trump: https://t.co/S56QbW5K5C
4	"The police are the most mistreated people in this country," Donald Trump #BlueLivesMatter #Trump2016 https://t.co/WfJvWUkMaB	Donald Trump says wages are too high. Really? Hardworking Americans don't think so. https://t.co/5oEK9UhGI1 https://t.co/1z0tuCedJa
5	Nobody beats me on National Security. https://t.co/sCrj4Ha1I5	I've released 9 years of tax returns. RT if you agree its time for Donald Trump to release his! https://t.co/08whtFVC0r

relatively better community separation. Because there were many tweets with positive sentiment that were correlated to pro-Jamala, Umigon also performed better than it did in the other cases. Table II shows the top 5 tweets from each polarity from the separation achieved using our algorithm. Note that the tweets on the left column are congratulating the winner (pro-Jamala), sharing winning related news, or talking against Russia. On the other hand, the tweets on the right column are against Jamala, and pointing out reasons for the deportation of Crimean Tatars, or arguing that the winning song should be disqualified. Some of the tweets are also sarcastic.

C. Trump

Donald Trump is the Republican Party nominee for President of the United States in the 2016 election. There have been much debate and controversies around the candidate. Tweets were collected using a single keyword `Donald Trump`, during April 2016. Collected tweets show support by the pro-Trump polarity and the negative opinions or mockery posted by the anti-Trump polarity. For the purpose of generating the ROC curves, the largest 1000 assertions were manually annotated. There were 372 pro-Trump, 522 anti-Trump, and 106 neutral assertions.

Figure 6c compares quality of factorization with other baselines. In this particular scenario, performance of our algorithm is around 2% better than community detection. This is because the corresponding source-assertion network had

strong community separation, with only 10% nonpolarized assertions being lightly connected. EM-Social also performs reasonably to find the separations because of the same reason. Table III shows the top 5 tweets from each polarity from the separation achieved using our algorithm. Note that the tweets on the left column are strongly pro-Trump in nature and describing support for him or praising him. On the other hand, tweets on the right column are sharing the negative information about the candidate, and pointing out the controversies.

V. RELATED WORK

Presence of polarization in social networks has been studied in various contexts. Conover et al. [1] study *retweet*-based social networks and *mention*-based social networks in political contexts related to U.S. congressional elections. Guerra et al. [17] study polarization metrics for social networks. They argue that *modularity* is not directly applicable as a measure of polarity because even without polarization modular communities are present. Uncovering polarization in social networks is important in various contexts. Bakshy et al. [5] study polarization in the context of Facebook. Amin et al. [6], Kase et al. [18] study crowd-sensing and fact-finders in the context of war and conflict situations. In this paper, we solve the orthogonal problem of separating the polarity classes.

Polarization in social network can be viewed as a community detection or graph partitioning problem [19], [20]. We do not directly apply such techniques because of the presence

of neutral sources and assertions. Moreover, the requirement of fusing muliple signals to converge to an expected solution required an optimization framework. Sentiment analysis [9], [11] can also be viewed as a related technique to uncover polarization. However, in our case, sentiment analysis is not directly applicable because the positive and negative classes in sentiment analysis can be orthogonal to the polarization group in question. Moreover, sentiment analysis is a *supervised* technique, while our technique is *unsupervised*. Sentiment analysis can require training and language model to map the sense of the text to sentiments. Even after training, such techniques can miss the assertions that are composed of sentiment-neutral wording, but semantically biased toward a certain side. On the other hand, our method looks at the source information and exploits the network structure to uncover polarity. As it does not consider text information, assertions that are not well connected in the network can be misclassified.

Finding a social-influence network, or source-dependency network has been studied in prior literature [4], [21], [22]. In this paper, we use the maximum likelihood approach proposed by Netrapalli and Sanghavi [4] to generate the social dependency matrix used as an input to our algorithm. In addition, bagging [23] and boosting [24] are two main solutions in ensemble learning [25]. In this paper, we follow this idea by filtering out bad separations by identifying Jaccard distance among the candidates and bagging filtered candidates.

VI. Conclusions

In this paper, we have presented a matrix factorization and ensemble based gradient descent algorithm to uncover polarization in social networks. We have evaluated our algorithm in the context of ongoing disputes, conflicts, or controversies as polarized situations. Experiments show that it can separate the tweets of different polarities by looking just at the source-assertion network and the social dependency network, and can be more than 90% accurate. Our algorithm performs much better than supervised techniques like sentiment analysis. Moreover, it also performs around $20\% - 30\%$ better than the community detection approaches, when the separation between the sources or the assertions of different polarities is obscured because of the presence of a large neutral network. If a particular source or assertion is not well connected to the network, the method can misclassify. Correctly estimating such cases with the help of additional information, deriving confidence bounds for the detected polarity, and jointly estimating polarity of the tweet with its veracity will be addressed in future works.

Acknowledgments

References

[1] M. Conover, J. Ratkiewicz, M. Francisco, B. Goncalves, F. Menczer, and A. Flammini, "Political polarization on twitter," in *Proc. International AAAI Conference on Weblogs and Social Media*, 2011.

[2] D. Wang, M. T. Amin, S. Li, T. Abdelzaher, L. Kaplan, S. Gu, C. Pan, H. Liu, C. Aggarwal, R. Ganti, X. Wang, P. Mohapatra, B. Szymanski, and H. Le, "Humans as sensors: An estimation theoretic perspective," in *ACM/IEEE Conf. on Information Processing in Sensor Networks*, 2014.

[3] M. T. A. Amin, S. Li, M. R. Rahman, P. T. Seetharamu, S. Wang, T. Abdelzaher, I. Gupta, M. Srivatsa, R. Ganti, R. Ahmed, and H. Le, "SocialTrove: A self-summarizing storage service for social sensing," in *International Conference on Autonomic Computing (ICAC'15)*, 2015.

[4] P. Netrapalli and S. Sanghavi, "Learning the graph of epidemic cascades," in *Proc. 12th ACM Joint Intl. Conf. on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 2012.

[5] E. Bakshy, S. Messing, and L. A. Adamic, "Exposure to ideologically diverse news and opinion on facebook," *Science*, vol. 348, no. 6239, pp. 1130–1132, 2015.

[6] M. T. A. Amin, T. Abdelzaher, D. Wang, and B. Szymanski, "Crowd-sensing with polarized sources," in *Proc. 2014 IEEE Intl. Conference on Distributed Computing in Sensor Systems*, 2014, pp. 67–74.

[7] S.-H. Cha, "Comprehensive survey on distance/similarity measures between probability density functions," *City*, vol. 1, no. 2, p. 1, 2007.

[8] M. Steinbach, G. Karypis, and V. Kumar, "A comparison of document clustering techniques," in *Proceedings of Workshop on Text Mining, 6th ACM SIGKDD International Conference on Data Mining*, 2000.

[9] R. Socher, A. Perelygin, J. Y. Wu, J. Chuang, C. D. Manning, A. Y. Ng, and C. Potts, "Recursive deep models for semantic compositionality over a sentiment treebank," in *Proc. EMNLP*, vol. 1631, 2013.

[10] S. Kiritchenko, X. Zhu, and S. M. Mohammad, "Sentiment analysis of short informal texts," *J. Artif. Int. Res.*, vol. 50, no. 1, May 2014.

[11] C. Levallois, "Umigon: sentiment analysis on tweets based on terms lists and heuristics," in *7th Intl Workshop on Semantic Evaluation*, 2013.

[12] (2016, Jul) Sentiment140 - A Twitter Sentiment Analysis Tool. [Online]. Available: http://http://www.sentiment140.com//

[13] G. Karypis and V. Kumar, "A fast and high quality multilevel scheme for partitioning irregular graphs," *SIAM Journal on scientific Computing*, vol. 20, no. 1, pp. 359–392, 1998.

[14] S. Yao, M. T. Amin, L. Su, S. Hu, S. Li, S. Wang, Y. Zhao, T. Abdelzaher, L. Kaplan, C. Aggarwal, and A. Yener, "Recursive ground truth estimator for social data streams," in *Proc. IPSN*, 2016.

[15] M. Uddin, M. Amin, H. Le, T. Abdelzaher, B. Szymanski, and T. Nguyen, "On diversifying source selection in social sensing," in *9th International Conference on Networked Sensing Systems (INSS)*, 2012.

[16] D. Wang, M. Al Amin, T. Abdelzaher, D. Roth, C. Voss, L. Kaplan, S. Tratz, J. Laoudi, and D. Briesch, "Provenance-assisted classification in social networks," *IEEE J-STSP*, vol. 8, no. 4, pp. 624–637, 2014.

[17] P. H. C. Guerra, W. Meira Jr, C. Cardie, and R. Kleinberg, "A measure of polarization on social media networks based on community boundaries." in *Proc. Intl. AAAI Conference on Weblogs and Social Media*, 2013.

[18] S. E. Kase, E. K. Bowman, M. T. Amin, and T. Abdelzaher, "Exploiting social media for army operations: Syrian civil war use case," in *Proc. SPIE Defense, Security, and Sensing*, 2014.

[19] G.-J. Qi, C. C. Aggarwal, and T. S. Huang, "Online community detection in social sensing," in *Proc. 6th ACM International Conference on Web Search and Data Mining (WSDM)*, 2013, pp. 617–626.

[20] G.-J. Qi, C. Aggarwal, and T. Huang, "Community detection with edge content in social media networks," in *Data Engineering (ICDE), 2012 IEEE 28th International Conference on*, April 2012, pp. 534–545.

[21] D. Kempe, J. Kleinberg, and E. Tardos, "Maximizing the spread of influence through a social network," in *Proc. 9th ACM SIGKDD Intl. Conference on Knowledge Discovery and Data Mining*, 2003.

[22] S. A. Myers and J. Leskovec., "On the convexity of latent social network inference," in *Proc. Neural Information Processing Systems*, 2010.

[23] L. Breiman, "Bagging predictors," *Machine learning*, vol. 24, no. 2, pp. 123–140, 1996.

[24] R. E. Schapire and Y. Freund, *Boosting: Foundations and algorithms.* MIT press, 2012.

[25] Z.-H. Zhou, *Ensemble methods: foundations and algorithms.* CRC press, 2012.