

On the Complexity of Opinions and Online Discussions

Utkarsh Upadhyay
MPI for Software Systems
utkarshu@mpi-sws.org

Aasish Pappu
Yahoo! Research
aasishkp@oath.com

Abir De
MPI for Software Systems
ade@mpi-sws.org

Manuel Gomez-Rodriguez
MPI for Software Systems
manuelgr@mpi-sws.org

ABSTRACT

In an increasingly polarized world, demagogues who reduce complexity down to simple arguments based on emotion are gaining in popularity. Are opinions and online discussions falling into demagoguery? In this work, we aim to provide computational tools to investigate this question and, by doing so, explore the nature and complexity of online discussions and their space of opinions, uncovering where each participant lies.

More specifically, we present a modeling framework to construct latent representations of opinions in online discussions which are consistent with human judgments, as measured by online voting. If two opinions are close in the resulting latent space of opinions, it is because humans think they are similar. Our framework is theoretically grounded and establishes a surprising connection between opinion and voting models and the sign-rank of matrices. Moreover, it also provides a set of practical algorithms to both estimate the dimensionality of the latent space of opinions and infer where opinions expressed by the participants of an online discussion lie in this space. Experiments on a large dataset from Yahoo! News, Yahoo! Finance, Yahoo! Sports, and the Newsroom app show that many discussions are multisided, reveal a positive correlation between the complexity of a discussion, its linguistic diversity and its level of controversy, and show that our framework may be able to circumvent language nuances such as sarcasm or humor by relying on human judgments instead of textual analysis.

ACM Reference Format:

Utkarsh Upadhyay, Abir De, Aasish Pappu, and Manuel Gomez-Rodriguez. 2019. On the Complexity of Opinions and Online Discussions. In *The Twelfth ACM International Conference on Web Search and Data Mining (WSDM '19)*, February 11–15, 2019, Melbourne, VIC, Australia. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3289600.3290965>

1 INTRODUCTION

People join online discussions to, on the one hand, express their own opinions and, on the other hand, approve and disapprove the opinions expressed by others. In this context, there is a wide variety of online platforms that enable their users to approve or disapprove

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM '19, February 11–15, 2019, Melbourne, VIC, Australia

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5940-5/19/02...\$15.00

<https://doi.org/10.1145/3289600.3290965>

each others' comments *explicitly* using, e.g., upvotes and downvotes. Here, whenever a user upvotes or downvotes a comment in an online discussion, she reveals the relative position of her opinion with respect to the opinion expressed in the comment in a latent space of opinions. By leveraging this observation from multiple comments, upvotes and downvotes, our goal is to investigate the nature and complexity of an online discussion and its space of opinions, uncovering where each participant lies.

There is a long history of theoretical models [3, 17, 18, 38, 39] and empirical studies [4, 10–12, 15, 16, 27, 31] of opinions. However, most of this previous work has reduced (potentially) complex opinions down to real-valued numbers—they have assumed that opinions lie on the real line. While a unidimensional space of opinions may be sufficient to coarsely characterize people into 2 groups, e.g., left leaning vs right leaning or liberals vs conservatives, they may be lacking at accurately representing complex, multisided opinions in an online discussion. In fact, if such opinions exist, a unidimensional representation of opinions will be unable to explain many common voting patterns under some of the most popular voting models, as we will show in Section 2. That being said, there are two notable exceptions—aspect-oriented sentiment analysis [9, 25] and collaborative filtering based on matrix factorization [6, 22, 24]. The former relies heavily on ad hoc methods and textual analysis for determining the sides in online reviews rather than on human judgments as we do. The latter aims to predict users' tastes (or ratings) about a set of items (e.g., movies) based on a partial observation of user-items ratings. While it typically considers multiple sides (or factors), it differs from our work in several key aspects. First, users and items typically lie in two different latent spaces while, in our work, comments and voters lie in the same latent space. Second, it determines the number of dimensions (sides) of the latent spaces empirically (e.g., using cross validation). This is in contrast with our work, which determines the dimension from first principles. Finally, most approaches consider real-valued ratings while we consider (binary) votes.

Current work. Given an online discussion consisting of a set of comments, which are upvoted and downvoted by a set of voters, we first introduce a latent multidimensional representation of the opinions expressed in the comments and the opinions held by the voters. Then, we propose two voting models, one deterministic and another probabilistic, which leverage the above multidimensional representation to characterize the voting patterns within an online discussion. Under this characterization, it becomes apparent that the dimensionality of the latent space of opinions is a measure of complexity of the online discussion—along how many different axis




Figure 1: Our modeling framework. From left to right, given a (toy) online discussion with a set of comments C and voters \mathcal{V} , our framework maps the upvotes and downvotes into a partially observed sign matrix S . Within S , each row corresponds to a comment and each column corresponds to a voter. Each $+1$ entry indicates that the voter upvoted the comment, -1 indicates that she downvoted the comment, and $?$ indicates that the voter did not vote. Then, the framework represents the opinions expressed in the comments and those held by the voters as r -dimensional real-valued vectors lying in the same latent space of opinions. Finally, it provides a set of practical algorithms to both estimate the dimension r of the latent space of opinions as well as infer the vectors of opinions which are consistent with the partially observed sign matrix S .

can the opinions expressed in the comments and the opinions held by the voters differ. The representation of opinions in this latent space of opinions has a remarkable property: if two opinions are close (far away) in the latent space it is because the voters—the crowd—think that they are *similar* (*dissimilar*). Such a property may not hold for other representations of the opinions, e.g., those based only on the textual data in the comments [23, 30] because of nuances in the use of language. Motivated by these observations, we develop¹:

- (i) A polynomial time algorithm to determine an upper bound on the minimum dimensionality that a latent space of opinions needs to have so that they are able to explain a particular voting pattern under the deterministic voting model.
- (ii) An inference method based on quantifier elimination to recover the latent opinions from the observed voting patterns under the deterministic voting model.
- (iii) An inference method based on maximum likelihood estimation to recover the latent opinions from the observed voting patterns under the probabilistic voting model.

Finally, we experiment with a large dataset from Yahoo! News, Yahoo! Finance, Yahoo! Sports, and the Newsroom app, which consists of one day of online discussions about a wide variety of topics. Our analysis yields several interesting insights. We find that only $\sim 25\%$ of the online discussions we analyzed can be explained using a unidimensional representation of opinions, $\sim 60\%$ of them require a two dimensional representation, and the remaining ones require a greater number of dimensions. This provides empirical evidence that, to provide opinions representations that are coherent with human judgments, it may be often necessary to move beyond one dimension. The presence of multisided opinions is an indication that the discussion may not be falling prey to demagoguery. Such finding is also supported by a positive correlation between the dimension of a discussion and its linguistic diversity. Moreover, the estimated r -dimensional opinions allow us to predict upvotes/downvotes in a discussion more accurately than a state of the art matrix factorization method [26] and a logistic regression classifier [29]. In this context, whenever an online discussion can be represented using one dimensional opinions, the deterministic model achieves higher predictive performance than the probabilistic

model, which allows for noisy voting, provides more accurate predictions. This suggests that, whenever humans face more complex discussions, their judgments become less *predictable*. Moreover, we find that a positive correlation between the complexity of the discussions and the level of agreement among comments. Lastly, by looking at particular examples of online discussions, we show that our modeling framework, by relying on human judgments, may be able to circumvent language nuances like sarcasm and humor, which are often difficult to detect using natural language processing. The examples will also illustrate how the dimensions uncover the different *sides* of the discussion.

2 MODELING OPINIONS AND VOTES

At the very outset, the underlying mechanism behind voting on online discussions is fairly commonplace and straight-forward. Every time a user expresses an opinion by posting a new comment in an online discussion, other users can upvote (downvote) the comment to indicate that they agree (disagree) with the expressed opinion.

In this context, whenever a user upvotes or downvotes a comment, she reveals the relative position of her opinion with respect to the opinion expressed in the comment. By leveraging this observation to multiple comments, upvotes and downvotes, our modeling framework will be able to infer the relative positioning of comments in an online discussion, as judged by the crowd. Moreover, by doing so, it will also find a meaningful joint latent representation for the opinions expressed in an online discussion as well as the opinions held by the users who voted. In the remainder of the section, we formally introduce our modeling framework, starting from the data it is designed for.

Online voting data. We observe an online discussion consisting of a set of comments C which are upvoted and downvoted by a set of voters \mathcal{V} . Here, we keep track of *who* voted *what* by means of the variables $y_{ij} = \{\uparrow, \downarrow, \circ\}$, which indicate that voter $j \in \mathcal{V}$ upvoted, downvoted, or did not vote on comment $i \in C$, respectively. Then, we define a (partially) observed sign matrix $S = [s_{ij}]$, where each (i, j) -th entry is given by

$$s_{ij} = \begin{cases} +1 & \text{if } y_{ij} = \uparrow \\ -1 & \text{if } y_{ij} = \downarrow \\ ? & \text{if } y_{ij} = \circ, \end{cases} \quad (1)$$

¹Implementation at <https://github.com/Networks-Learning/discussion-complexity>.

the sign ? indicates that the voter did not vote and thus we cannot know whether she agrees (or disagrees) with the comment. We denote by Ω the set of indexes where we have observations, *i.e.*, $\Omega = \{(i, j) \mid s_{ij} \neq ?\}$. Figure 1 illustrates the above definitions for a given toy example.

Next, we introduce our multidimensional representation of the opinions expressed in the comments and those held by the voters and then elaborate on our voting model, which relates these opinions to the observed voting data.

Opinion representation. Unidimensional (scalar) real-valued representations of opinions, owing largely to their interpretability, have been used most commonly in the literature, following the example set by the seminal works by DeGroot [14] and Rowley [34]. Thus, we could think of using such unidimensional representation of opinions in our work. However, under that choice, we would be unable to explain certain voting patterns illustrated below, which are common in many online discussions.

Given an online discussion, assume we represent the opinions expressed in each comment $i \in C$ as $c_i \in \mathbb{R}$ and the opinion held by voter $j \in \mathcal{V}$ as $v_j \in \mathbb{R}$. Now, we elaborate separately on two of the most popular voting models in the literature [28]: the proximity model and the directional model. Under the proximity model, the voters use the Euclidean distance as a similarity measure and decide to cast an upvote if $|v_j - c_i| \leq \theta$, where θ is a threshold, and a downvote otherwise. Now consider the voting pattern 1 in Figure 2a. It is easy to show that there are no real-valued scalar opinions v_1, v_2, v_3 and c_1, c_2, c_3 leading to such a voting pattern: assume that $v_1 \leq v_2 \leq v_3$ (as the pattern is symmetric, we can always relabel the voters and comments to make this true) and $c_2 < v_2$. Then $|v_2 - c_2| > \theta \implies v_2 > c_2 + \theta$, and $|v_3 - c_2| \leq \theta \implies v_3 \leq c_2 + \theta$. This contradicts the assumption that $v_2 \leq v_3$. We arrive at a similar contradiction with the assumption $c_2 > v_2$.

Under the directional model, the voters use the dot product as a similarity measure and decide to cast an upvote if $v_j \cdot c_i \geq 0$ and a downvote otherwise. Here, consider the voting pattern 2 in Figure 2b. Again, it is easy to show that there are no real-valued non-zero scalar opinions v_1, v_2, v_3 and c_1, c_2, c_3 leading to such a voting pattern. The first row requires $v_1 \cdot c_1 \geq 0$ and $v_2 \cdot c_1 < 0$, which implies $\text{sign}(v_1) \neq \text{sign}(v_2)$. However, the second row requires $v_1 \cdot c_2 \geq 0$ and $v_2 \cdot c_2 \geq 0$, which implies $\text{sign}(v_1) = \text{sign}(v_2)$ and this leads to a contradiction.

Motivated by the above examples, given an online discussion, we represent the opinions expressed in the comments and those held by the voters as r -dimensional real-valued vectors lying in the same latent space. More formally, we represent the opinion expressed in each comment $i \in C$ as $c_i \in \mathbb{R}^r$ and we stack all these opinions into a matrix C , in which the i -th row corresponds to the opinion c_i^T . Similarly, we represent the opinions held by each voter $j \in \mathcal{V}$ as $v_j \in \mathbb{R}^r$ and stack all these opinions into a matrix V , in which the j -th row corresponds to the opinion v_j^T . Here, one can think of the dimension r as a measure of the complexity of the online discussion—along how many different *axis* can the opinions expressed in the comments and the opinions held by the voters differ. Figure 1 illustrates the above definitions using a toy example.

		← \mathcal{V} →			
		1	2	3	
↑	1	↓	↑	↑	
↷	2	↑	↓	↑	
↓	3	↑	↑	↓	

(a) Voting pattern 1

		← \mathcal{V} →			
		1	2	3	
↑	1	↑	↑	↑	
↷	2	↑	↓	↑	
↓	3	↑	↑	↑	

(b) Voting pattern 2

Figure 2: Examples of unfeasible voting patterns under the proximity and directional voting models with unidimensional (scalar) real-valued representation of opinions.

Voting model. Given a comment i which expresses an opinion c_i and a voter j who holds an opinion v_j , we introduce two voting models, one deterministic and another probabilistic, inspired by the directional model of voting discussed above.

— *Deterministic voting model:* In this model, we can uniquely determine each vote y_{ij} from the comment’s opinion c_i and voter’s opinion v_j by means of the following deterministic rule:

$$y_{ij} = \begin{cases} \uparrow & \text{if } \langle c_i, v_j \rangle \geq 0 \\ \downarrow & \text{if } \langle c_i, v_j \rangle < 0. \end{cases} \quad (2)$$

In the above rule, the vote y_{ij} depends on the *angle* between the opinion vectors c_i and v_j —if the angle is greater (less) than 90° , *i.e.*, c_i and v_j lie in the same (different) half-plane in the latent space, then $y_{ij} = \uparrow$ (\downarrow).

Under this voting model and a partially observed sign-matrix S derived from votes using Eq. 1, two natural question emerge:

- (i) What is the minimum dimension r of the latent space needed to recover the observed entries in S from the above decision rule without errors?
- (ii) Once we know the minimum dimension r , can we infer the opinion vectors c_i and v_j ?

We will answer both questions affirmatively in Section 3 and 4, respectively.

— *Probabilistic voting model:* In the definition of our deterministic model, we have implicitly assumed that voters do not make any *errors* while casting their votes. However, this assumption might be rather restrictive in some scenarios. To overcome this, we also propose a probabilistic voting model in which votes are binary random variables Y_{ij} , and,

$$\mathbb{P}[Y_{ij} = y_{ij}] = p(y_{ij}) = \frac{1}{1 + \exp(-s_{ij} \langle c_i, v_j \rangle)}, \quad (3)$$

where $s_{ij} = +1$ if $y_{ij} = \uparrow$ and $s_{ij} = -1$ if $y_{ij} = \downarrow$. Similarly, as in the case of the deterministic model, we will propose a method to infer the opinion vectors c_i and v_j under this model in Section 4. In doing so, we will make the assumption that all the latent opinions are finite, *i.e.*, $\exists \alpha > 0. \|C\|_\infty \leq \alpha \wedge \|V\|_\infty \leq \alpha$.

Remark. In the above model definitions, we opt for a similarity metric based on dot products because the euclidean distance, used in the proximity model, does not *scale* well with increasing dimensionality: the relative *volume* of the opinion space where a voter will cast an upvote is proportional to $(\theta/\alpha)^r$ where θ is the threshold for the user, r is the dimension of the latent space and α is the upper bound on the opinion values.




Figure 3: Illustration of Case 1 and Case 2 while selecting the edges of the minimum spanning tree which will help determine the permutation of the rows of the matrix that minimizes $SC^*(S)$. First, the edge (c_1, c_2) is selected and constraint $s_{13} \equiv s_{23}$ is added to \mathcal{E} . Next, the edge (c_1, c_3) is selected and s_{13} and s_{23} are filled with $s_{33} = -$.

3 COMPLEXITY OF ONLINE DISCUSSIONS

In this section, we present an algorithm which can determine an upper bound on the minimum dimension r that a latent space of opinions needs to have so that C and V are able to *explain* voting patterns exhibited by the voters which result in a particular vote-matrix S under the deterministic voting model, i.e., $\forall (i, j) \in \Omega$, $s_{ij} = \text{sign}(m_{ij})$, where $[m_{ij}] = CV^T$. To this aim, we will first introduce the notion of sign-rank of a sign matrix. Then, we will show that the problem of determining r reduces to finding the sign-rank of a partially observed sign matrix. Finally, we will present an efficient algorithm to estimate the sign-rank.

Sign-rank of a sign matrix. Paturi and Simon [32] introduced the classical notion of sign-rank of a sign matrix, which is closely related to the VC dimension of concept classes [2], as follows:

DEFINITION 1. Let M be a real matrix and $\text{sign}(M)$ denote a matrix such that $\forall i, j$, $(\text{sign}(M))_{ij} = \text{sign}(M_{ij})$. Then, the sign-rank of a sign matrix S is defined as:

$$\text{sign-rank}(S) = \min \{ \text{rank}(M) \mid \text{sign}(M) = S \}.$$

Here, we extend the above definition to partially observed sign matrices as follows:

DEFINITION 2. The sign rank of a partially observed sign matrix S is defined as:

$$\text{sign-rank}(S) = \min \{ \text{rank}(M) \mid \forall (i, j) \in \Omega, \text{sign}(M)_{ij} = s_{ij} \}.$$

It is easy to see that, if the rank of a matrix M is r , then we can decompose the matrix into two components of the form CV^T using, e.g., the singular value decomposition (SVD). Hence, the problem of determining r reduces to the problem of finding $\text{sign-rank}(S)$.

Note that the sign-rank of a sign matrix can be much lower than its actual rank, as was noticed by Hsieh et al. [19] in the context of signed graph models. For example, consider the sign-rank of the matrix $B = 2\mathbb{I}_n - 1_n$, where \mathbb{I}_n is the identity matrix and 1_n is the matrix of all 1 of size $n \times n$. For $n \geq 3$, $\text{sign-rank}(B)$ remains 3 though the matrix itself is always of full rank n . Moreover, note that, in our setting, the sign-rank does not merely correspond to the *number* of topics being discussed in an online discussion. Instead, the complexity may be manifest in the combination of the topics under discussion: the voters may agree with some opinions in a comment while disagreeing with others.

Estimating the sign-rank of a partially observed sign matrix.

The problem of determining whether $\text{sign-rank}(S)$ is 1 can be solved by a simple breath-first search (BFS). We first create a signed bipartite graph of comments and voters with adjacency matrix S . Then for each connected component in the graph, pick one $(i, j) \in \Omega$, set

$c_i = +1$ and $v_j = s_{ij}$, and fill in the remaining values using BFS by multiplying the source node value with the sign of the edge to arrive at the destination node value. The intuition is that if voter j has down (up) voted comment i , then i and j have opposite (same) polarity. If a consistent assignment of ± 1 to all the nodes is possible, then $\text{sign-rank}(S) = 1$.

However, this algorithm does not generalize to multiple dimensions. To estimate the sign-rank of a partially observed sign matrix, we adapt the algorithm for (fully observed) sign matrices proposed recently by Alon et al. [2]. First, we explain the main ideas behind the original algorithm and then describe the necessary, non trivial modifications we propose.

The original algorithm upper-bounds the sign-rank of a (fully observed) sign matrix S by the number of *sign-changes* in the columns of the matrix. More formally, define the function $SC(S)$ as the maximum number of sign changes in any column of the matrix S , i.e., $SC(S) = \max_j |\{i \mid s_{i,j} \neq s_{(i+1),j}\}|$, $\text{Sym}(S)$ as the set of all possible row permutations of S , and the function $SC^*(S)$ as:

$$SC^*(S) = \min_{S' \in \text{Sym}(S)} SC(S').$$

Then, the following lemma establishes the relationship between the sign-rank of matrix S and $SC^*(S)$, which the original algorithm exploits [1]:

LEMMA 3. For a sign matrix S , $\text{sign-rank}(S) \leq SC^*(S) + 1$.

To use the above result, our algorithm needs to do two tasks:

- Find a matrix $\bar{S} = [\bar{s}_{ij}]$ such that it is a *completion* of S , i.e.,

$$\bar{s}'_{ij} = \begin{cases} s_{ij} & \text{if } (i, j) \in \Omega \\ \pm 1 & \text{if } (i, j) \notin \Omega. \end{cases} \quad (4)$$

- Find $S' \in \text{Sym}(\bar{S})$ such that it minimizes the maximum number of sign-changes in its columns.

The algorithm will output $SC(S') + 1$ as the estimated sign-rank.

Our algorithm does both tasks together while it computes an estimation of $SC^*(S)$ using an algorithm by Welzl [37, See Ex. 4]. In a nutshell, we construct a graph in which each node corresponds to a row of the partially observed matrix S and the weight of each edge between nodes u and v is given by the number of columns where the signs of the corresponding rows disagree. Then, we extract a spanning tree from the completely connected graph which minimizes the number of sign-changes between pairs of vertices connected by an edge, as shown in Algorithm 1. In the process of creating the spanning tree, we also fill the matrix. Finally, we derive a permutation of the rows based on the tree to construct S' .

More in detail, to understand how Algorithm 1 fills the matrix as it computes the spanning tree, we distinguish two different cases:

Algorithm 1 Construct a sign-minimizing spanning tree for the columns of S .

Input: A $[|C|] \times [|\mathcal{V}|]$ sign-matrix $S = [s_{ij}]$
Output: Spanning-tree of the rows.
1: $l \leftarrow 0$; $F \leftarrow \{\}$; $w : \mathbb{Z}^2 \rightarrow \mathbb{R}$; $Z \leftarrow [|C|]$; $Y \leftarrow [|\mathcal{V}|]$
2: **while** $l < |C|$ **do**
3: **for** $u \in Z$, $v \in [|C|]$; $(u, v) \cup F$ is cycle-free **do**
4: $w(\{u, v\}) += |\{j \in Y \mid s_{uj} \neq s_{vj} \wedge s_{uj} \neq ? \wedge s_{vj} \neq ?\}|$
5: **end for**
6: $\{u^*, v^*\} \leftarrow \operatorname{argmin} w(\cdot)$; $S, Y, Z \leftarrow \text{UPDATE}(\{u^*, v^*\}, S)$
7: $F \leftarrow F \cup \{u^*, v^*\}$; $l \leftarrow l + 1$;
8: **end while**
9: **return** F

Case 1: When, given a column, only one of the rows has a missing entry. Consider, for example, we have two rows $u = (+1, +1, ?, +1)$ and $v = (-1, +1, -1, -1)$. To calculate the weight of the edge between u and v , i.e., $w(\{u, v\})$ in line 4 of Algorithm 1, we ignore the second column, as $s_{u,2} = s_{v,2}$, and the third column, as it has a missing entry $s_{u,3} = ?$. Hence, we report $w(\{u, v\}) = 2$, because signs of u and v differ in 1st and 4th column. Now, if this edge was chosen in line 6 of Algorithm 1, we *modify* u , such that this weight indeed is the true weight of the edge: we replace $s_{u,3} = ?$ with the corresponding value in v , i.e., $s_{v,3} = -1$, via line 5 in Algorithm 2.

Case 2: When, given a column, both entries are unknown. If $u = (+1, +1, ?, +1)$ and $v = (+1, -1, ?, +1)$, we would still calculate the weight the same way as above. Hence, $w(\{u, v\}) = 1$. However, if this edge was chosen in line 6 of Algorithm 1, we could keep the weight the same by merely ensuring that both u and v have the *same* value in the third column, i.e., $s_{u,3} \equiv s_{v,3}$. Hence, we create and save the constraint that the third column of u and v must always have the same value in line 9 of Algorithm 2. Now, say a few steps into the creation of the spanning tree, we find that the missing value in u has to be set to -1 as it was being picked as part of an edge under Case 1 above. Then, we can also set the same value in the third column of v , i.e., $s_{v,3} \leftarrow +1$, via line 5 in Algorithm 2. Note that since each column in S contains at least 1 entry which is ± 1 , as each voter has voted at least once, we will eventually hit Case 1 and fill in all missing entries. This process is illustrated in Figure 3, where the first step creates an equivalence $s_{13} \equiv s_{23}$, and the second step fills in the missing entries with s_{33} .

Note that we are *conservative* and greedy while filling in the missing entries, i.e., the edge selected at the l th step will have the same *weight* $w(e)$ at the l th iteration if the algorithm was to be run on the filled matrix and it will be the minimum weight amongst all valid edges at step l (though the edge may not be unique in having that weight). Additionally, our algorithm ensures that the weight of the edge selected at iteration l is the minimum possible, given the history of selections. After obtaining a spanning tree, one can *walk* the tree by performing a depth-first search starting from any source node and create a permutation of the rows by dropping the duplicate nodes in the walk. Hence, we can obtain $S' \in \text{Sym}(S)$ and report $SC(S') + 1$ as r .

Then, we can establish the upper bound on the dimension by using the following series of self-evident inequalities: $\text{sign-rank}(S) \leq \text{sign-rank}(S') \leq SC^*(S') + 1 \leq SC(S') + 1$.

Algorithm 2 UPDATE procedure used in Algorithm 1.

Operation $\leftarrow \frac{\mathcal{E}}{Y, Z}$ assigns RHS to all entries equivalent to the LHS in \mathcal{E} (line 9) and updates Y, Z .

Input: $\{u, v\}$ edge chosen; **and** $S = [s_{ij}]$ sign-matrix;
Output: Updated S and all rows/columns with updated signs.
1: // \mathcal{E} initialized as an empty set and persisted across calls.
2: $Z \leftarrow \{\}$; $Y \leftarrow \{\}$
3: **for** $j \in [|\mathcal{V}|]$ **do**
4: **if** $s_{uj} = ?$ **and** $s_{vj} \neq ?$ **then**
5: $s_{uj} \leftarrow \frac{\mathcal{E}}{Y, Z} s_{vj}$
6: **else if** $s_{uj} \neq ?$ **and** $s_{vj} = ?$ **then**
7: $s_{vj} \leftarrow \frac{\mathcal{E}}{Y, Z} s_{uj}$
8: **else if** $s_{uj} = ?$ **and** $s_{vj} = ?$ **then**
9: $\mathcal{E} \leftarrow \mathcal{E} \cup \{s_{uj} \equiv s_{vj}\}$
10: **end if**
11: **end for**
12: **return** S, Y, Z

Finally, we would like to highlight that the spanning tree algorithm presented above minimizes the *average* number of sign-changes in S' . Welzl [37] also describe a variant of the algorithm which produces guarantees on the worst case number of sign-changes in S' ; the way the *weight* $w(\cdot)$ is calculated is more involved in the variant. This variant was used by Alon et al. [2] to design the first polynomial time algorithm with approximation guarantees for the sign-rank of the matrix S' . Remarkably, the UPDATE procedure in Algorithm 2 can be ported to that variant without any changes, to complete a partially observed matrix S matrix with worst case guarantees as well. However, that version is computationally more expensive, more complex, and does not offer significantly better results in practice in our dataset. Hence, for ease of exposition, we have described the simpler of the two versions.

Computational complexity. The computational complexity of Algorithm 1 can be determined by the computations needed for each missing element in the matrix S . Except on the first initialization iteration ($l = 0$), the loop on line 3 will be executed once for each missing element (u, i) in the initial S , right after s_{ui} is fixed via UPDATE. Hence, the work done for each missing entry, in the worst case, will be incrementing $w(\cdot)$ by 1 $\forall v \in [|\mathcal{V}|]$ such as $s_{uj} \neq s_{vj}$ in the loop on line 3. This can be done in $O(|C| \log |C|)$ time if $w(\cdot)$ is implemented as a priority queue. As there are at most $|C| \times |\mathcal{V}|$ missing entries, the computational cost of Algorithm 1 is $O(|C|^2 |\mathcal{V}| \log |C|)$. All other operations, i.e., initialization of $w(\cdot)$, calculating argmin , checking for cycles in the tree, the creation of an walk, maintaining \mathcal{E} , etc., have lower complexity.

Remark. In our implementation, we do a (non-exhaustive) search over walks with different sources to improve our estimate of r and break ties in calculating the $\operatorname{argmin} w(\cdot)$ randomly. Also, as $\text{sign-rank}(S) = \text{sign-rank}(S^T)$, we run the algorithm on both matrices and report the smaller value.

4 MULTISIDED OPINION ESTIMATION

Given an online discussion, we infer the corresponding r -dimensional opinions C and V for the deterministic and probabilistic voting models introduced in Section 2 as follows.




Figure 4: Distribution of number of observed elements and fraction of unique voting patterns in matrix S .

Deterministic voting model. By definition, under the deterministic voting model, we know that the corresponding r -dimensional opinions V and C and the partially observed sign matrix S need to satisfy the following inequalities:

$$\bigwedge_{(i,j) \in \Omega} s_{ij} \left(\sum_{k=1}^r c_{ik} \cdot v_{jk} \right) > 0 \quad (5)$$

where c_{ik} and v_{jk} are the k -th entry of the opinions c_i and v_j , respectively, and Ω is the set of observed entries in S . However, we also know from Section 3 that, for each voting pattern, there will be a minimum dimension r_{min} under which such an opinion embedding will not exist.

This reduces the problem of finding the opinions V and C to the existential theory of reals [36] and, for small values of r and moderate number of comments and voters, $|C|$ and $|V|$, this problem can be solved via a procedure called quantifier elimination [20], using, e.g., the solver Z3 [13]. This procedure eliminates the inequalities in the disjunction in Eq. 5 one by one by discovering subsets of $\mathbb{R}^{r \times (|V|+|C|)}$ where they are satisfied, backtracking to select a different region when an inequality cannot be satisfied, and using sound heuristics to prune the search. The procedure terminates when *any* assignment to each variable is found or when all regions of $\mathbb{R}^{r \times (|V|+|C|)}$ are eliminated.

Note that, if $r < r_{min}$, the solver will conclude that the problem is unsatisfiable. Hence, by iteratively increasing r and checking for satisfiability of Eq. 5, one could determine the true sign-rank of any matrix S . However, as the most efficient method known for quantifier elimination is doubly exponential in the number of variables, calculating the minimum dimension r in this way would be computationally more expensive than using the polynomial algorithm introduced in Section 3.

Probabilistic voting model. Given a partially observed matrix S , under the probabilistic voting model, we estimate V and C by solving the following constrained maximum likelihood estimation problem with hyperparameters α :

$$\begin{aligned} & \underset{C, V}{\text{maximize}} && - \sum_{(i,j) \in \Omega} \log \left(1 + \exp \left(- s_{ij} \langle c_i, v_j \rangle \right) \right) \\ & \text{subject to} && \|C\|_{\infty} \leq \alpha, \|V\|_{\infty} \leq \alpha. \end{aligned} \quad (6)$$

The structure of the above problem allows us to adapt an efficient 1-bit matrix completion method based on stochastic gradient descent [8]. Finally, note that unlike in the deterministic model, for each voting pattern and dimension r , there will always exist opinions C and V that best fit the data.

Dim.	Discuss.	$ C $	$ V $	Patterns
1*	1,139	19.7 ± 16.7	19.9 ± 20.1	16.2 ± 14.8
2*	2,820	51.5 ± 85.4	57.6 ± 75.2	46.0 ± 57.3
3	97	43.7 ± 35.3	56.0 ± 34.5	45.4 ± 24.6
4	88	148 ± 147	195 ± 194	149 ± 125
5	245	247 ± 272	296 ± 267	218 ± 179
6	126	445 ± 461	470 ± 366	354 ± 275
7	89	598 ± 555	706 ± 518	512 ± 355
≥ 8	112	2596 ± 2867	2785 ± 2540	1831 ± 1602

Table 1: Number of comments, voters and unique voting patterns seen in the dataset for discussions with different dimensions. The numbers in each column are the mean values \pm the standard deviation. Dimensions marked with * indicate that they were determined using Z3 and are the true dimensions of the discussions. While the dimension of discussions is positively correlated with the size and participants, discussions of different complexity can be found on the entire spectrum, as is also shown in Figure 5.




Figure 5: Distributions of the number of comments $|C|$ per discussion for different dimension values (Dim). The distributions of the number of voters and voting patterns exhibit similar behavior.


Remark. In both models, the estimated opinions are unique up-to orthogonal transformations since the inequalities in Eq. 5 and the likelihood in Eq. 6 only depend on entries of CV^T and $(CO)(VO)^T = C(OO^T)V^T = CV^T$ for any orthogonal matrix O .

5 EXPERIMENTS

Data description. Our dataset contains $\sim 19,800$ online discussions, each associated to an article from Yahoo! News (including contributed articles), Yahoo! Finance, Yahoo! Sports, and the Newsroom app, which contain ~ 5 million votes, cast by $\sim 200,000$ voters on $\sim 685,000$ comments, posted by $\sim 151,000$ users. These votes were randomly sampled from all votes which were cast on comments made by users in the US on August 8, 2017.

As a pre-processing step, we discard discussions with less than 10 comments, as they contain too little data to provide meaningful results. After this step, our dataset consists of 4,700 discussions, with ~ 4.5 million votes, cast by $\sim 199,000$ voters, on $\sim 645,000$ comments, posted by $\sim 137,000$ users. Figure 4 shows the *richness* of the data in the votes gathered in the online discussions by means of the sparsity of S and the number of unique columns of S , which we name as *voting patterns*.

Complexity of discussions. In this section, we compute the complexity of the discussions, *i.e.*, the dimensionality of the latent space of opinions, for the online discussions in our dataset. For each online discussion, we determine whether it can be explained using an unidimensional space of opinions using the linear time algorithm presented at the beginning of Section 3. If it cannot be explained



(a) Discussions with dimension 2 (b) Discussions with dimension 3

Figure 6: Performance of our dimensionality estimation algorithm.





Figure 7: Lexical similarity of an online discussion against its dimension. The lexical similarity is the mean Jaccard similarity of the lexical tokens used in all pairs of comments in the discussion. The discussions are classified by the number of comments as small (smallest 33%), medium, and large (largest 33%) and data is shown only for dimensions which contain more than 5 discussions.

using one dimension, we determine whether it can be explained using a two- or three-dimensional space of opinions via quantifier elimination², following Section 4. Finally, if it cannot be explained using two or three dimension, we resort to the algorithm presented in Section 3, which provides an upper bound on the true dimension.

Table 1 summarizes the results, which show that the opinions of about 1,139 (25%) of the discussions can be explained using one dimension, 2,820 (60%) of the discussions require two dimensions, while the remaining 757 discussions (15%) require a higher number of dimensions. This allows us to conclude that the opinions in most of the online daily discussions (85%) can be explained using a latent space of relatively low dimensions, *i.e.*, $r \leq 3$. Moreover, while discussions with a higher number of participants ($|C| + |V|$) and richness (*i.e.*, higher number of voting patterns and lower sparsity) require, in general, a latent space of opinions with a larger number of dimensions, there is a large variability spanning the entire spectrum of online discussions, as shown in Figure 5. Next, we evaluate how tight is the upper bound on the true dimension provided by our algorithm for online discussions, which we used above for discussions whose dimension we could not find using quantifier elimination. To this aim, we run our algorithm on discussions whose true dimension we could find using quantifier elimination and compare the upper bound with the true dimension. Figure 6 summarizes the results, which show that, for discussions whose opinions can be explained using two (three) dimensions, our algorithm recovers the true dimension for 48% (46%) of the discussions and is off by one for 37% (44%) of them.


Finally, we investigate the relationship between the complexity of the discussions, estimated using human judgments, and their linguistic diversity, estimated using textual features. To this aim, for

²In practice, we found quantifier elimination to be sufficiently scalable to test whether an online discussion can be explained using up to two dimensions.



(a) Discussions with dimension 1 (b) Discussions with dimension 2

Figure 8: Vote prediction accuracy for the deterministic voting model (DVM), the probabilistic voting model (PVM), a state of the art matrix factorization method [26] (MF), and a logistic regression classifier [29] (LR) using textual features extracted using Rake [33]. Moreover, the performance for DVM, PVM and MR uniformly increases as the number of unique voting patterns increases, in contrast, the performance for LR remains relatively constant.




(a) Mutual agreement and upvotes (b) Agreement/upvote distribution

Figure 9: Agreement and percentage of upvotes among all votes in discussions. Agreement is measured in terms of percentage of comment pairs (c_i, c_j) for which $c_i^T c_j > 0$. The dashed lines show the values of the metrics if the up/down votes or the latent dimensions were randomly distributed.

each online discussion, we compute the average Jaccard similarity of the lexical tokens used in all pairs of comments as a measure of lexical similarity. Figure 7 summarizes the results, which show a positive correlation between the complexity of a discussion and its linguistic diversity, as one may have expected.

Opinions in online discussions. In this section, we first evaluate both quantitatively and qualitatively the *quality* of the estimated r -dimensional opinions in the online discussions and then leverage the estimated opinions to shed some light on the level of controversy in online discussions. Here, we used the opinion estimation method for the probabilistic voting model introduced in Section 4, which scales gracefully with the dimension r .


In terms of quantitative evaluation, we assess to which extent the deterministic voting model (DVM) and the probabilistic voting model (PVM) can predict whether a voter will upvote or downvote a comment from the estimated opinions in comparison with two baseline methods: (i) a state of the art matrix factorization method [26] (MF), which assumes the entries in S are real valued, and (ii) a logistic regression classifier [29] (LR) that uses 200,000 keywords extracted using Rake [33] as features. To this aim, for each discussion, we held out some of the observed upvotes and downvotes, estimate the opinions from the remaining votes, and then predict the votes from the held-out set. However, since our data is very sparse, as shown in Figure 4, and even holding out a small fraction of votes may change the underlying dimension of the latent space of opinions, we resort to leave-one-out validation.



C_0 :	[...] [Donald Trump] has [...] Enquirers ^a [which] he considers a treasure trove of information.
C_1 :	He should change his name to Donald J Dubious.
C_2 :	[...] Trump can be an #\$\$%, and Islam can be cancer [...] they are not mutually exclusive [...]
C_3 :	Why not? Try anything. Terrorism has got to stop now!
C_4 :	It is a great idea
C_5 :	Trump family motto-“It’s not a lie if you believe it.”

^aNational Enquirers is a well known entertainment magazine in US.

Figure 10: A subset of comments, estimated multidimensional opinions and unidimensional sentiments for an online discussion about politics. Two pairs of comments, i.e., (C_0, C_1) and (C_3, C_4) , express a similar opinion, however, the lexical overlap between comments within each pair is low. By leveraging the judgments of the voters, our method is able to identify they are similar, and their estimated opinions lie close to each other in the latent space of opinions, however, sentiment analysis is unable to do so and the unidimensional sentiments do not lie close to each other. Moreover, the estimated opinion of a comment expressing an opposite view to the ones above, C_2 , lies in an orthogonal direction.



C_0 :	You can forget about \$16 for a while.
C_1 :	Bye-Bye twitter sweet 16.
C_2 :	[...] when world leaders speak, they turn to Twitter first. [...] How is it trading at \$16? [...] How come Dorsey ^a can’t monetize this instantaneous platform?
C_3 :	[...] It’s about time [for] more positive news to get it [...] up again. [...] Seems to have support \$16ish. [...]
C_4 :	[Wall Street/CNBC] only want to pump [selective] stocks [...] Twitter of China, Weibo, is selling for \$88.00 a share [...]

^aJack Dorsey is the CEO of Twitter.

Figure 11: A subset of comments, estimated multidimensional opinions and unidimensional sentiments for an online discussion about finance. There are two distinct issues being discussed: (i) the price of a stock (C_0, C_1, C_2, C_3) and (ii) a discussion about reasons for the price (C_2, C_4) . C_0 and C_1 say humorously that the price will stay below \$16, while C_2 and C_3 suggest that the price *may* rise up. C_4 suggests Wall Street/media bias against the stock and is neutral about its price, while C_2 questions the management of the company.

Moreover, we randomly select 200 discussions to tune the hyperparameters of PVM and these discussions are excluded from the validation set. LR was regularized using 10-fold cross-validation and MF was provided the true rank of the underlying matrix as input. Figure 8 summarizes the results, which show that:

- (i) DVM (PVM) beats all other methods for discussions with dimension 1 (2).
- (ii) The performance of DVM, PVM and MF increases as the number of unique voting patterns in the dataset increase, in contrast, the performance of LR, which uses text features, does not benefit much from additional voting patterns.
- (iii) While for discussions where opinions can be explained using two dimensions, PVM achieves better performance, for discussions which require only one dimension, DVM beats PVM. A potential explanation for this behavior is that, whenever humans face simpler decisions, i.e., their opinions can be explained using one dimension, they become more *predictable*.

In terms of qualitative evaluation, we first assess to which extent comments in online discussions agree (or disagree) by analyzing the estimated opinion embeddings of the comments. More specifically, for each online discussion, we compute the percentage of distinct comment pairs (c_i, c_j) for which $c_i^T c_j > 0$ and compare this quantity with the percentage of upvotes among all votes (upvotes and downvotes). Figure 9a summarizes the results, which show that

the higher the dimensionality of the latent space of opinions, the lower the agreement between comments, as one may have expected. Remarkably, such finding would not be apparent directly from the relatively constant fraction of upvotes. However, relative upvotes are typically the measure of *controversy* (or, rather, consensus) employed by various websites, like Reddit, to sort articles/comments.

Finally, we take a close look into the comments, inferred multidimensional opinions and unidimensional sentiment³ of a discussion about politics, shown in Figure 10, and a discussion about finance, shown in Figure 11. The discussion about politics shows that, even if the lexical overlap between comments which express a similar opinion is low, e.g., C_0 and C_2 or C_4 and C_5 , our opinion estimation method is able to identify they are similar, as a human would do, by leveraging the judgments of the voters. Note that, due to their low lexical overlap, it would be difficult to identify such similarity using methods based on textual analysis, as revealed by the unidimensional sentiments. The discussion about the price of Twitter stock (see Figure 11) shows that our method is able to capture objective opinions about the price (whether it stays at \$16 or goes up, C_3 , or stays down, C_0 and C_1), along one axis and subjective opinions questioning the reason behind the price drop along a different axis (suggesting management is the reason, C_2 , or media bias/corruption

³Comment sentiments were calculated using Convolutional Neural Networks trained on Stanford Sentiment Treebank [21, 35].

in Wall Street, C_4). Note that C_2 suggests both, that the price should go up, and that the reason for the decrease is the management. Also in this case, an analysis of the unidimensional sentiments would not reveal these rich relationships among the comments.

6 CONCLUSION

In this work, we have proposed a modeling framework to generate latent representations of opinions using human judgments, as measured by online voting. As a consequence, such representations exhibit a remarkable semantic property: if two opinions are close in the latent space of opinions, it is because the voters—the crowd—think that they are similar. Our modeling framework is theoretically grounded and establishes an unexplored, surprising connection between opinion and voting models and the sign-rank of a matrix. Moreover, it also provides a set of practical algorithms to both estimate the dimension of the latent space of opinions and infer where opinions expressed in comments and held by voters lie in this space. Experiments on a large dataset from Yahoo! News show that many discussions are multisided and avoid falling prey to demagoguery, provide insights into human judgments and opinions, and show that our framework is able to circumvent language nuances, e.g., sarcasm and humor, by relying on human judgments.

Our work also opens up many interesting venues for future work. For example, our measure of complexity—the dimension of the latent space of opinions—may be a good starting point to develop theoretically grounded measures of polarization [10, 11, 27] and controversy [15, 16], which have been lacking in the literature. Moreover, it would be very interesting to augment our modeling framework to also incorporate, in addition to the voting data, the textual information in the comments, the identity of commenters, and their trust-worthiness. Besides increasing the accuracy of our method, these may aid the interpretability of the dimensions as well. Our algorithm for determining the minimum dimension under which the opinion space is able to explain the voting data exhibits weak theoretical guarantees though it performs well on real-data. It would be interesting to develop exact algorithm by adapting recent advances in exact sign-rank estimation [5, 7].

Acknowledgements. We thank Mounia Lalmas, Dmitry Chistikov, and Rupak Majumdar for useful discussions.

REFERENCES

- [1] Noga Alon, Peter Frankl, and V. Rödl. 1985. Geometrical realization of set systems and probabilistic communication complexity. In *FOCS*.
- [2] Noga Alon, Shay Moran, and Amir Yehudayoff. 2016. Sign rank versus VC dimension. In *Conference on Learning Theory*.
- [3] R. Axelrod. 1997. The dissemination of culture: a model with local convergence and global polarization. *Journal of conflict resolution* 41, 2 (1997), 203–226.
- [4] Pablo Barberá. 2015. Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data. *Political Analysis* 23, 1 (2015), 76–91.
- [5] Ronen Basri, Pedro F Felzenszwalb, Ross B Girshick, David W Jacobs, and Caroline J Klivans. 2009. Visibility constraints on features of 3D objects. In *CVPR*.
- [6] James Bennett, Stan Lanning, et al. 2007. The Netflix prize. In *Proceedings of KDD cup and workshop*. Vol. 2007. New York, NY, USA, 35.
- [7] Amey Bhangale and Swastik Kopparty. 2015. The complexity of computing the minimum rank of a sign pattern matrix. *arXiv preprint arXiv:1503.04486* (2015).
- [8] Sonia A Bhaskar and Adel Javanmard. 2015. 1-bit matrix completion under exact low-rank constraint. In *CISS*.
- [9] Jürgen Broß. 2013. Aspect-oriented sentiment analysis of customer reviews using distant supervision techniques. *Doctoral dissertation* (2013).
- [10] Yoonjung Choi, Yuchul Jung, and Sung-Hyon Myaeng. 2010. Identifying controversial issues and their sub-topics in news articles. In *Pacific-Asia Workshop on Intelligence and Security Informatics*. Springer, 140–153.
- [11] Michael Conover, Jacob Ratkiewicz, Matthew R Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political polarization on Twitter. *ICWSM* (2011).
- [12] Abir De, Sourangshu Bhattacharya, Parantapa Bhattacharya, Niloy Ganguly, and Soumen Chakrabarti. 2014. Learning a linear influence model from transient opinion dynamics. In *CIKM*.
- [13] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An efficient SMT solver. *Tools and Algorithms for the Construction and Analysis of Systems* (2008), 337–340.
- [14] Morris H DeGroot. 1974. Reaching a consensus. *J. Amer. Statist. Assoc.* 69, 345 (1974), 118–121.
- [15] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Quantifying Controversy on Social Media. *ACM Transactions on Social Computing* 1, 1 (2018), 3.
- [16] Pedro Henrique Calais Guerra, Wagner Meira Jr, Claire Cardie, and Robert Kleinberg. 2013. A Measure of Polarization on Social Media Networks Based on Community Boundaries. In *ICWSM*.
- [17] R. Hegselmann and U. Krause. 2002. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation* 5, 3 (2002).
- [18] P. Holme and M. E. Newman. 2006. Nonequilibrium phase transition in the coevolution of networks and opinions. *Physical Review E* 74, 5 (2006), 056108.
- [19] Cho-Jui Hsieh, Kai-Yang Chiang, and Inderjit S Dhillon. 2012. Low rank modeling of signed networks. In *KDD*.
- [20] Dejan Jovanović and Leonardo de Moura. 2012. Solving Non-linear Arithmetic. In *Automated Reasoning - 6th International Joint Conference, IJCAR 2012, Manchester, UK, June 26-29, 2012. Proceedings (Lecture Notes in Computer Science)*, Vol. 7364. Springer, 339–354.
- [21] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1746–1751.
- [22] Yehuda Koren and Robert Bell. 2015. Advances in collaborative filtering. In *Recommender systems handbook*. Springer, 77–118.
- [23] Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *International Conference on Machine Learning*. 1188–1196.
- [24] Greg Linden, Brent Smith, and Jeremy York. 2003. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet computing* 1 (2003), 76–80.
- [25] Bing Liu. 2012. Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies* 5, 1 (2012), 1–167.
- [26] Rahul Mazumder, Trevor Hastie, and Robert Tibshirani. 2010. Spectral regularization algorithms for learning large incomplete matrices. *Journal of machine learning research* 11, Aug (2010), 2287–2322.
- [27] Yelena Mejova, Amy X Zhang, Nicholas Diakopoulos, and Carlos Castillo. 2014. Controversy and sentiment in online news. *arXiv preprint arXiv:1409.8152* (2014).
- [28] Samuel Merrill and Bernard Grofman. 1999. *A unified theory of voting: Directional and proximity spatial models*. Cambridge University Press.
- [29] Kevin Murphy. 2012. *Machine learning: a probabilistic perspective*. MIT Press, Cambridge, UK.
- [30] Hamid Palangi, Li Deng, Yelong Shen, Jianfeng Gao, Xiaodong He, Jianshu Chen, Xinying Song, and Rabab Ward. 2016. Deep sentence embedding using long short-term memory networks: Analysis and application to information retrieval. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)* 24, 4 (2016), 694–707.
- [31] Bo Pang, Lillian Lee, et al. 2008. Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval* 2, 1–2 (2008), 1–135.
- [32] Ramamohan Paturi and Janos Simon. 1984. Probabilistic communication complexity. In *FOCS*.
- [33] Stuart Rose, Dave Engel, Nick Cramer, and Wendy Cowley. 2010. Automatic keyword extraction from individual documents. *Text Mining: Applications and Theory* (2010), 1–20.
- [34] Charles K Rowley. 1984. The relevance of the median voter theorem. *Zeitschrift für die gesamte Staatswissenschaft/Journal of Institutional and Theoretical Economics* H. 1 (1984), 104–126.
- [35] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*. 1631–1642.
- [36] Alfred Tarski. 1998. A decision method for elementary algebra and geometry. In *Quantifier elimination and cylindrical algebraic decomposition*. Springer, 24–84.
- [37] Emo Welzl. 1988. Partition trees for triangle counting and other range searching problems. In *Proceedings of the fourth annual symposium on Computational geometry*. ACM, 23–33.
- [38] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione. 2013. Binary opinion dynamics with stubborn agents. *ACM Transactions on Economics and Computation* 1, 4 (2013), 19.
- [39] M. E. Yildiz, R. Pagliari, A. Ozdaglar, and A. Scaglione. 2010. Voting models in random networks. In *Information Theory and Applications Workshop*. 1–7.