Voting with Social Networks: Truth Springs from Argument Amongst Friends

John A. Doucette, Hadi Hosseini, Alan Tsang, Kate Larson, and Robin Cohen
David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
{j3doucet,h5hossei,akhtsang,klarson,rcohen}@cs.uwaterloo.ca

ABSTRACT

In certain voting problems, a central authority must infer a hidden ground truth by aggregating the opinions of an electorate. When individual assessments are drawn i.i.d. and are correct with probability $p > 0.5$, aggregating enough votes will yield the ground truth with high probability. However, in reality voters' opinions are often influenced by those of their friends. In certain social networks, this interaction may cause the aggregated opinions to be misleading. A center may use the network structure to recover the original distribution of votes from the post-interaction opinions of voters, and so recover the ground truth in spite of confounding discussions. In this paper, we consider a series of novel models of these underlying social interactions, and derive maximum likelihood estimators for the ground truth in these models, given the social network and votes. We also evaluate these new estimators, as well as existing ones, on common classes of social networks, and derive analytical bounds on the estimators' performance in different environments. In all, we provide important insights into admitting influence from peers during voting.

1. INTRODUCTION

The study of voting is conventionally divided into two distinct, but related, perspectives. In one view, voting is the process of arriving at some compromise given the subjective opinions of the voters. In a political contest for instance, voters' different interests are expressed as preferences, and aggregated to select a winner that represents the will of the electorate. Voters may have deep-seated philosophical differences, and tend to (subjectively) view various outcomes in distinct ways because of these differences. In this view of voting, there is no presumption on the 'correctness' of any particular view, and instead winners are ultimately selected for voters on the basis of a rule that allows a collective compromise.

In contrast, we can view voting as the process of aggregating individuals' beliefs regarding an objective truth about the world. This view starts from the assumption that voters are casting ballots representing their (potentially noisy or erroneous) impressions of some event. For example, imagine a sporting event where viewers perceive the field from many different angles (as in a circular stadium). When a ball comes close to a goal line, different viewers will have different perceptions about whether the ball has crossed the line. Despite their distinct viewpoints, there is a *true* answer (either the ball really did cross the line, or it really did not). By having the viewers vote on whether the ball crossed, we could aggregate their perspectives on what has really happened, and recover the underlying truth. Early in the history of social choice, Condorcet formalized this notion with the *Jury Theorem*, stating that, provided voters have a better than even chance of observing the truth (rather than observing the incorrect answer), aggregating enough votes will yield the truth with high probability [3].

Appears at: 2nd Workshop on Exploring Beyond the Worst Case in Computational Social Choice. Held as part of the 14th International Conference on Autonomous Agents and Multiagent Systems. May 4th, 2015. Istanbul, Turkey.

This paper focuses on the second view, which has applications to many problems in artificial intelligence, and to multiagent systems in particular (e.g. multiagent resource allocation, aggregating noisy sensor data from a cooperative swarm of agents, recommender systems).

To date, most existing work in the domain of voting with a ground truth has assumed that voters' individual impressions of the truth are generated *independently*. In practice, however, this may not be so. In recent work, Conitzer [8] considers the possibility that voters' impressions of the truth might be influenced by discussions with other members of the electorate, raising the question "Should social network structure be taken into account in elections?". This additional factor can confound attempts to recover the truth by aggregating individual opinions. As a trivial example, consider the social network depicted in Figure 1.

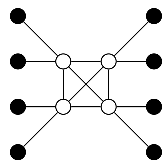

Figure 1: An example of voting on a social network.

Each node represents a voter in this social network. A central 4-clique have the belief that the correct answer is 'white', while peripheral neighbors connected only to the members of the clique believe the correct answer is 'black'. Suppose the true answer is black, and, consistent with the requirements of the Jury Theorem, nodes are more likely (by a 2:1 ratio) to have observed the truth. However, if social interactions occur according to this graph, the members of the clique could reinforce each others' opinions, while the isolated perimeter nodes would perceive only opinions opposed to their own, and consequently might change their minds, and cast a vote for 'white' based on the local aggregation of information from their neighbors. Naively aggregating the views of the voters after social interactions of this kind take place (e.g. on the basis of majority opinion) could, thus, result in an overwhelming majority for white, when in fact, the true answer is black. This suggests that aggregation in the presence of a social network should incorporate information about the network structure into its decisions.

In Conitzer's earlier work [7, 8], two preliminary models for social interaction were considered, but neither was entirely satisfactory. In this paper, we present a new model for recovering the ground truth under the assumption that *it is easier to convince someone of the truth than of a falsehood*, which we call the "righteous argument" model. We derive a maximum likelihood estimator for the truth given the social network and the observed votes after social interactions, under this assumption, which can be computed efficiently. The model is extended to the case of hierarchical social networks, and to the case where outcomes are non-binary. The model is evaluated using simulations, and compared to a naive model that does not take the structure of the social network into account. We also provide a theoretical assessment of the model, and describe the class of problems where it will provide an advantage over counting the votes without heed for the structure of the social network. We conclude with a discussion of the model and directions for future work, and by answering Conitzer's question "Should social network structure be taken into account in elections?" in the affirmative, at least for some elections.

2. RELATED WORK

Social networks represent a model of how people communicate, interact and influence each other in their daily lives. Properly harnessed, a social network could be used to promote the spread of certain ideas, and curtail others. This idea motivated much of early research in the area. The field of innovation diffusion investigates the adoption of new technologies by members of a community [6, 15]. The core premise of innovation diffusion is the notion of peer imitation: once a number of peers have 'vetted' and adopted a new technology, one is incentivized to adopt the technology as well, or risk being left behind.

More recent work on diffusion in social networks focuses on the diffusion of opinions. At its core, the idea is that as people interact with their peers, their opinions (modeled as a continuous value) will become more similar. Early mathematicians explored linear models of such opinion dynamics, which admitted closed form solutions [11]. Kraus was the first to investigate nonlinear systems, bringing in agent based modeling techniques from the artificial intelligence community to simulate networks that were hard to examine analytically [12]. Since then, AI researchers have become increasingly interested in social networks. Subsequent works explore opinion dynamics using random graph models such as scale-free networks which mimic key characteristics of real world networks [2]. More recent works incorporate elements of social choice in social network settings. An early work by Martins combines the continuous 'opinion space' with discrete expressions reminiscent of voting [13]. Gasper et al. [10] investigate voting behavior in the presence of social influence; Tsang and Larson explore the effects of skepticism between agents [16].

Our work follows in a similar vein of examining social choice mechanisms in the presence of a social network. Our approach differs most markedly from that of opinion dynamics in that we are not examining long term behavior of the network; rather, we emulate the time critical aspects of real world problems, and conduct only a brief round of interactions before measuring opinions from the network.

Other researchers have also examined the design of social choice mechanisms for a social network. Boldi et al. [4] propose a system of delegative democracy that operates on a social network. Here, the goal is to create and analyze a concrete voting system for implementing representative democracy, where each agent may delegate another to vote for them by proxy. This delegation process is transitive, with the weight of the delegated ballot diminishing exponentially with distance. This work differs from ours in that it proposes a mechanism that allows agents to correlate their votes (by allowing another to vote in their stead), whereas our system assumes that opinion correlation will occur among neighbors in a social network, and examines the impact this has on how voting may be used to uncover a hidden ground truth.

Salehi-Abari and Boutilier [14] also examine the voting behavior of agents within a social network. Their motivating example is a group of friends selecting a restaurant together. While people may have subjective preferences over the restaurants, the enjoyment experienced by their friends will play a crucial role. They propose an efficient mechanism by which social welfare may be maximized while accounting for these empathic utilities. Thus, Salehi-Abari and Boutilier examine the other side of the social network coin, where utilities depend not only on one's own preferences (and the outcome), but also those of one's friends.

The work most closely related to our own is Conitzer's recent work proposing a maximum likelihood approach to voting over a social network [8]. Conitzer considers the problem of estimating the winner given the network structure and the votes cast after some (unspecified) social interactions have taken place on the network. The initial model considered by Conitzer [7] was shown to reduce to a naive counting of votes, without incorporating network structure at all. Recognizing this, Conitzer [8] proposed an alternative model where instead of opinions being distributed over the vertices of the graph (i.e. the voters), opinions were distributed over the *edges* of the graph (i.e. the conversations or interactions between voters). In this model, the assumption is that individual conversations are more likely to produce the correct answer than an incorrect one, and that voters cast their ballot according to the majority of the opinions present in their interactions. The center (the entity that is aggregating the votes) still observes only the final ballots cast, and does not directly observe the edges. The MLE for this model is somewhat complex, because there are many possible assignments of opinions to the edges that could have generated a given vote profile. Finding the correct winner under this model was shown to be $\#P$-hard, but an efficient estimator exists for a reasonable variation of the model.

Although this model provides an interesting view, it seems counter-intuitive to assign opinions to edges (conversations) rather than vertices (voters). Presumably the opinions that arise from a given conversation follow from the information (and thus opinions) of the participants, and are not an intrinsic property of the conversation itself. Additionally, the fact that the opinions of the conversations are distributed independent of the participants creates peculiar situations where voters who both ultimately agree (and vote) that the answer is correct, might actually have a conversation where they agree that the answer is a different, incorrect, alternative. These drawbacks motivate the creation of a new model.

3. A MODEL FOR SOCIAL INTERACTIONS

Let $G = \langle V, E \rangle$ denote an undirected graph representation of a social structure where V is the set of all vertices where $|V| = n$, and E is the set of all edges between the vertices. In our social structure, each voter is represented by a vertex $v \in V$, and each edge $(u, v) \in E$ represents an edge between the vertices u and v.

Given a set of alternatives C, the goal of a voting problem in the presence of social network is to select a winner $c \in C$, according to a voting rule, that represents the will of the voters in the underlying social network. For simplicity, we first assume binary alternatives, i.e., $C = \{c, c'\}$.[1] Each voter casts a vote $A_v \in A_V$. The collective votes of voter v's neighbors is $A_{Nei(v)}$, and is simply defined as $A_{Nei(v)} = \{A_u | (u, v) \in E\}$. The probability of a voter casting a ballot for a candidate, denoted by f, is a function of its idiosyncratic belief about the correct outcome and its neighbors votes.

3.1 The Naive Model

We first review the model proposed by Conitzer [7] and discuss its drawbacks. In this model, the probability distribution of a voter's report is comprised from two independent parameters, the probability of observing a vote A_v given that c was the correct answer is denoted by $g(A_v|c)$, and the probability of observing the joint profile of votes for A_v and its neighbors $A_{Nei(v)}$ is represented by $h(A_v, A_{Nei(v)})$.

Given a social network, the probability of a voter v casting a particular ballot, along with the ballot profile of its neighbors, given that the correct answer is c, is given by

$$f(A_v, A_{Nei(v)}|c) = g(A_v|c)\, h(A_v, A_{Nei(v)})$$

Intuitively, g captures the tendency of voters to vote for the correct answer (like in the Jury Theorem), and h captures the tendency of voters to agree (or perhaps disagree) with their neighbors. If we then wished to estimate the likelihood of an observed ballot profile, under the assumption that a candidate c is the true winner, the likelihood is then proportionate to

$$\mathcal{L}(A_V|\hat{c} = c) \propto \prod_v f(A_v, A_{Nei(v)}|c)$$

under the assumption that the votes are conditionally independent of each other except with respect to their neighbors. The winner is the candidate $c \in C$ that maximizes the likelihood of the vote profile.

The main drawback of this model is the independence of the social interactions h from the outcome c, which results in identical values of h when computing the likelihood under each candidate. Thus, the social interaction value h can be ignored when maximizing the likelihood, and the model reduces to

$$\mathcal{L}(A_V|\hat{c} = c) \propto \prod_v g(A_v|c)$$

which does not incorporate information from the social network structure at all. Indeed, under the assumption that $g(A_v|c)$ is > 0.5 when c is the correct answer, and < 0.5 otherwise, this model represents Condorcet's Jury Theorem [3], and is maximized by picking the candidate with the most votes as the winner.

These shortcomings in particular arose from the fact that h, the joint probability of a voter's reported opinion and the opinions of its neighbors, did not depend on the outcome c. Therefore, we propose a new model based on the notion of righteous argument that takes the social interactions into account.

3.2 The Righteous Argument Model

A more realistic model of social interactions is to suppose that the social influence of neighbors in a network depends on the true outcome. The *righteous argument* model assumes that between two connected voters, the voter with the correct opinion about the truth is more likely to convince the voter with an incorrect opinion, than the other way around.

We assume that the social influence h is dependent on $\hat{c}$ (the true outcome), and that the social influence $h(A_v, A_{Nei(v)}|\hat{c})$ can be partitioned into a product of independent pairwise interactions. Let $h'(A_v, A_u|\hat{c} = c)$ be the probability that we observe vertices v and u, which are connected in the social network, cast votes A_v and A_u respectively, given that the correct answer is c. Then, the joint probability of observing a particular vote alongside those in the neighborhood of the voter that cast it is given by

$$h(A_v, A_{Nei(v)}|\hat{c} = c) = \prod_{u \in Nei(v)} h'(A_v, A_u|\hat{c} = c)$$

Given the pair-wise interaction function h', the Maximum Likelihood Estimator (MLE) for the observed vote profile can be written as

$$\mathcal{L}(A_V|\hat{c} = c) \propto \prod_v g(A_v|\hat{c} = c)\, h(A_v, A_{Nei(v)}|\hat{c} = c)$$

A winner is a candidate that satisfies the following

$$\arg\max_{c \in C} \mathcal{L}(A_V|\hat{c} = c)$$

For binary candidates, we assume $g(A_v|c) = p$ when $A_v = c$ and $\dot{p} = 1 - p$ otherwise. Thus, the likelihood function reduces to

$$\mathcal{L}(A_V|\hat{c} = c) \propto p^i \dot{p}^{n-i} \prod_V h(A_v, A_{Nei(v)}|\hat{c} = c)$$

where i is the total number of reported votes for c in A_V. Notice that, if two adjacent nodes u and v cast differing votes, then the terms $h'(A_v = c, A_u = c'|\hat{c} = c)$ and $h'(A_v = c, A_u = c'|\hat{c} = c)$ will appear in the likelihood for both candidates, and so such discordant edges need not be considered. This leaves only the concordant edges where both parties agreed. Suppose that $h'(A_v = c, A_u = c|\hat{c} = c) = q$ for some $q > 0.3$, and that $h'(A_v = c', A_u = c'|\hat{c} = c) = \dot{q}$ for some $\dot{q} < 0.3$[2]. Then we can write the final likelihood as

$$\mathcal{L}(A_V|\hat{c} = c) \propto p^i \dot{p}^{(n-i)} q^x \dot{q}^y$$

where x is the number of connections between nodes with opinion c, and y is the the number of connections between nodes with a different opinion.

Recall that p, q and $\dot{q}$ are actually unknown quantities. It follows that this model can produce a certain result only when one candidate has both a majority of votes and a majority of the concordant edges in the graph between vertices that voted for the winning candidate. Computing the most likely candidate is then linear in the total size of the social network.

Observe also, that although this model assumes the function h can be decomposed into h' (and so that the probabilities of observing particular correlations between votes is distributed independently), it will recover the truth on models which generate votes in a correlated fashion, provided that such models are *anonymous* (i.e. voters have the same prior probability of being connected to one another, and of having particular opinions). That is, voters' initial votes ($V_i \in V$) and connections are distributed independently prior to the realization of the social structure and votes. The model assumes that the probability $P(V_i = c|c)$, prior to knowledge of the particular graph structure or exact distribution of opinions among the voters, is symmetric across all voters. In a similar vein, prior to exact knowledge of the graph structure, one can compute the prior probability that $P(V_i = V_j = c|c)$, and the prior probability that any two voters i and j are connected in the graph. As long as these prior probabilities are symmetric for all pairs and individuals, then

[1] Section 3.4 provides a thorough consideration of multi-candidate voting.

[2] By the Righteous Argument assumption, $P(c, c|c) > P(c, c') = P(c', c) > P(c', c')$. It follows that $q \geq \frac{1}{3}$ and $\dot{q} < \frac{1}{3}$.

the assumptions of the model hold. Note that even if a process is much more likely to generate graphs with homophily, or to cause voters to adopt the correct opinion according to the joint distribution over their neighbors' views, the *prior* probability of such concordant pairs (before knowledge of the graph structure) will still be symmetric unless the method cares about the *names* of the voters. This allows for decomposition.

3.3 Asymmetric Social Influence

A natural extension to our righteous argument model with symmetric social influence (i.e. where when A influences B, B also influences A) is to consider social structures wherein individuals' have different social interactions with their neighbors in social networks such as boss-employee relationships, the reach of popular media, and unusually stubborn and steadfast individuals. In particular, pairwise interactions are not necessarily symmetric, that is, for each pair of nodes v and u in graph $G = \langle V, E\rangle$, $(u, v) \in E$ represents a directed edge from u to v, and does not imply $(v, u) \in E$.

Let the in-neighbors of v be $\overleftarrow{Nei(v)} = \{u \mid (u, v) \in E\}$ and the out-neighbors of v be $\overrightarrow{Nei(v)} = \{u \mid (v, u) \in E\}$. The sets of votes by in- and out- neighbors are represented by $A_{\overleftarrow{Nei(v)}}$ and $A_{\overrightarrow{Nei(v)}}$ respectively. The in-neighbors of v represent those individuals who are influenced by v's opinion, and the out-neighbors of v represent the individuals who influence v's opinion; these sets need not be disjoint.

Note that the Naive model [7] explained in Section 3.1 does not change under the asymmetric assumption as the social influence function h is independent of the outcome c, thus the network structure has no impact on maximizing the likelihood of the observed votes.

Incorporating the objective truth of voters' opinions into their ability to influence their in-neighbors, function $f(A_v, A_{Nei(v)}|\hat{c})$ is decomposed into three components: $g(A_v|\hat{c})$, $\overleftarrow{h}(A_v, A_{\overleftarrow{Nei(v)}}|\hat{c})$, and $\overrightarrow{h}(A_v, A_{\overrightarrow{Nei(v)}}|\hat{c})$.[3] Expanding the social influence model to independent pairwise factors, we have

$$\overleftarrow{h}(A_v, A_{\overleftarrow{Nei(v)}}|\hat{c}) = \prod_{A_u \in A_{\overleftarrow{Nei(v)}}} h'(A_u, A_v|\hat{c})$$

$$\overrightarrow{h}(A_v, A_{\overrightarrow{Nei(v)}}|\hat{c}) = \prod_{A_u \in A_{\overrightarrow{Nei(v)}}} h'(A_v, A_u|\hat{c})$$

Here, $h'(A_v, A_u|\hat{c})$ represents the probability that the configuration A_v and A_u was achieved after v attempted to influence u's opinion ($(u, v) \in E$), given the correct choice $\hat{c}$. Based on the pairwise interactions, we define the following distinct interaction cases:

I **Enlightened pair**: Let $q = h'(A_v = A_u = c|\hat{c} = c)$ be the probability that two vertices with an interaction agree on the correct outcome.

II **Unenlightened pair**: Let $\dot{q} = h'(A_v = A_u = c'|\hat{c} = c)$ be the probability that two vertices with an interaction agree on the incorrect outcome.

III **Successful Resistance**: Let $r = h'(A_v = c', A_u = c|\hat{c} = c)$ denote the probability that node u was unconvinced by the arguments of node v toward the incorrect alternative.

IV **Failed to Enlighten**: Let $\dot{r} = h'(A_v = c, A_j = c'|\hat{c} = c)$ denote the probability that node u was unconvinced by the arguments of v toward the correct alternative.

If we assume that interactions increase the likelihood of finding the correct outcome, then it follows that $q > \dot{q}$, and $r > \dot{r}$. We can calculate the probability of a voting profile, proportionate to:

$$\mathcal{L}(A_V|\hat{c}) \propto \prod_{v\in V} g(A_v|\hat{c}) \Big[\prod_{u\in\overleftarrow{Nei(v)}} h'(A_v, A_u|\hat{c}) \prod_{u\in\overrightarrow{Nei(v}} h'(A_u, A_v|\hat{c}) \Big]$$

And selecting the most likely winner is equivalent to:

$$\arg\max_{c\in C} \mathcal{L}(A_V|\hat{c} = c)$$

Then, let x be the number of votes for c, y be the number of edges where both end points voted for c (case I), $\bar{y}$ be the number of edges where both endpoints voted for c' (case II), and z and $\bar{z}$ be the number of edges that disagree according to case III and case IV respectively. Then, the likelihood for c is maximized as

$$p^x \dot{p}^{n-x} q^y \dot{q}^{\bar{y}} r^z \dot{r}^{\bar{z}}$$

Conversely, the likelihood for c' is

$$p^{n-x} \dot{p}^x q^{\bar{y}} \dot{q}^y r^{\bar{z}} \dot{r}^z$$

Therefore, c wins if and only the following holds true:

$$p^x \dot{p}^{n-x} q^y \dot{q}^{\bar{y}} r^z \dot{r}^{\bar{z}} > p^{n-x} \dot{p}^x q^{\bar{y}} \dot{q}^y r^{\bar{z}} \dot{r}^z$$

Which can be simplified to

$$p^{2x-n} q^{y-\bar{y}} r^{z-\bar{z}} > \dot{p}^{2x-n} \dot{q}^{y-\bar{y}} \dot{r}^{z-\bar{z}}$$

This implies that there is a clear solution for a winning outcome if $2x > n$, $y > \bar{y}$ and $z > \bar{z}$; or $2x < n$, $y < \bar{y}$ and $z < \bar{z}$. That is, a clear winner exists when one candidate has a majority of votes, a majority of both Case I over Case II, and Case III over Case IV edges. Again, this can be computed in time proportional to the size of the social network.

3.4 Multiple Candidates

In this section, we consider an extension to the binary voting model presented thus far when there are more than two candidates, where there is only one correct choice among those candidates.[4]

Let C be the set of alternatives under consideration. Define the probability that voter v votes for candidate A_v, given that candidate c is the correct alternative, as $f(A_v, A_{Nei(v)}|\hat{c} = c)$. For simplicity we assume that voters cast a ballot containing only a single candidate, and not a full rank ordering. As in the naive model, suppose that $f(A_v, A_{Nei(v)}|\hat{c} = c) = g(A_v|\hat{c} = c)\, h(A_v, A_{Nei(v)})$. Then, similar to the undirected case, the structure of the social network cannot influence the outcome. Assuming that voters in general will vote for the true outcome $c = \hat{c}$ more often than for any other, an MLE for this model is simply to count the votes as they appear. If we extend the model using our assumption of righteous argument, there are four possible cases for any pair of adjacent vertices, with unique interpretations:

I **Enlightened Pair:** Let $q = h'(A_v = A_u = c|\hat{c} = c)$ be the probability that two adjacent vertices agree on the correct outcome.

[3]Despite the state of a node being influenced only by its out-neighbors, we need to examine both $\overleftarrow{h}$ and $\overrightarrow{h}$, because we are evaluating likelihood of observing the configuration of the entire graph.

[4]Crucially, the ground truth does not specify a ranking of the candidates, but that one of them is intrinsically correct while the others are *equally* incorrect.

II **Failed to Enlighten:** Let $\dot{r} = h'(A_v = c, A_u = c_2|\hat{c} = c)$, where $c_2 \in C \setminus c$ be the probability that a vertex with the correct outcome is adjacent to a vertex with an incorrect outcome.

III **Unenlightened Pair:** Let $\dot{q} = h'(A_v = A_u = c_2|\hat{c} = c)$, where $c_2 \in C \setminus c$ be the probability that two adjacent vertices both vote for the same incorrect candidate.

IV **Unenlightened Argument:** Let $r = h'(A_v = c_2, A_u = c_3|\hat{c} = c)$, where $c_2, c_3 \in C \setminus c$ be the probability that two adjacent vertices both vote for a different incorrect candidates.

Finding the MLE for this model is equivalent to finding:

$$\arg\max_{\hat{c} \in C} p^x \dot{p}^{n-x} q^y \dot{q}^{\bar{y}} r^z \dot{r}^{\bar{z}}$$

which imposes constraints on declaring a winner similar to the directed case discussed above.

Note that this model conforms only to the case where there is a single candidate that is correct, and the other outcomes are, in some sense, equally wrong. If instead there exists a *ranking* over the outcomes, such that some outcomes are perhaps "more true" than others, the inference problem is substantially altered. A natural extension would be to assume an order on the probabilities associated with different pairs of candidates appearing in adjacent votes.

4. EVALUATION

We now evaluate the performance of the models described above on simulated votes held within various social networks. We consider only the case of binary outcomes, but evaluate across several families of random graphs, including both directed and undirected graphs, and a range of possible problem parameterizations.

To generate a problem instance, we generate a social network with n vertices (the voters), according to a particular (parameterized) graph generation algorithm. We then assign an initial opinion to each voter, with n_c of the n voters starting with the correct opinion, and the remainder receiving the incorrect opinion. A model of influence dynamics is then applied to the graph, so that opinions of voters tend to change to match those of their neighbours. The final product is a social network where each vertex is assigned a vote, and where the votes are the result of initial opinions modified by discussion between the voters.

4.1 Influence Dynamics

Let $Nei(v)$ denote those voters whose opinions have a direct influence on v when the latter updates. In undirected graphs, $Nei(v)$ includes exactly the neighbors of v and v itself. In directed graphs, $Nei(v)$ includes exactly the in-neighbors of v and v itself.

We consider a process by which voters revise their opinions one by one, with parameter k controlling the total number of voters whose opinions have changed. The process also incorporates the idea of righteous argument, where each discussion has a higher chance of swaying a voter toward the correct state than the incorrect state. Voters who hold an incorrect opinion are more likely to be swayed to the correct position when many of their friends already hold the correct opinion. On the other hand, a voter who already holds the correct opinion is unlikely to be swayed away from that position (independent of the number of friends with incorrect opinions).

This simulation models situations where voters are convinced only by arguments for the truth (or, with some small constant probability, may revert to a false viewpoint), and where the probability of an interaction with any given neighbor is constant and independent over time (so voters with more neighbors that vote correctly wait less time to talk to someone with the truth).

Formally, our influence dynamics iteratively picks k voters whose opinions will be changed. Let $P(v)$ denote the probability a voter v is selected:

$$P(v) = \begin{cases} \frac{1}{Z}, & \text{if } A_v = c \\ \frac{1}{Z}(|\{u|u \in Nei(v), A_u = c\}| + 1), & \text{if } A_v = c' \end{cases}$$

where Z is a normalization constant such that the probabilities over all voters sum to 1. These probabilities are recomputed after each voter opinion update, and so a voters' opinion may change multiple times over the course of a simulation.

4.2 Graph Models

We consider four different models of random graphs: the classic Erdös-Rényi undirected random graph (ER), the Barabási-Albert preferential attachment model (BA), and the natural directed counterparts of these models (dER and dBA, respectively).

The Erdös-Rényi random graph [9] is a model that incorporates a minimum number of assumptions. Given a connection probability pr, an Erdös-Rényi random graph is generated by connecting distinct vertices u, v with probability pr. The natural extension of this model to directed graphs adds the edge (u, v) with probability pr, which is independent from the addition of the reverse edge (v, u).

The Barabási-Albert random graph [1] is generated via preferential attachment, where new vertices are 'drawn' toward high-degree vertices. It leads to the generation of several high degree 'hub' nodes that characterize scale-free networks, and are frequently used to model human generated social networks. A Barabási-Albert random graph with attachment parameter m is generated by repeatedly adding vertices to a network, and connecting the new vertex with m existing vertices, each chosen with probability proportional to their respective degrees.

Our extension of the Barabási-Albert model to the directed case depicts a strongly hierarchical network, where newly added vertices will look towards influential older vertices as an example and a source of sound judgment. Each new vertex v is added with an edge to m existing vertices; an existing vertex u is chosen to an in-neighbor of v with probability proportional to $d_{in}(u) + \delta$.[5] This is a simplification of the directed scale-free graphs proposed by Bollobás [5], and produces directed acyclic graphs.

4.3 Aggregation

We compare two methods of aggregating opinions from nodes in a social network: (1) The naive model, where the aggregated opinion is the majority opinion from the votes at the vertices, and (2) the Righteous Arguments model, where we examine both the majority opinion from the vertices, as well as the distribution of opinions along the edges of the network.

In an undirected graph, we need only examine the opinions from concordant edges. An edge (u, v) in G is a *concordant edge* for c if $A_u = A_v = c$. If state c wins a majority from both the vertices and concordant edges, then the model has successfully determined the correct outcome; if c' wins a majority from both, the model has failed. If the majority outcomes are different, then the model is undecided: it cannot decide on an outcome without further assumptions about the underlying probabilities, and outputs "Undecided".

In a directed graph, we examine both the opinions from concordant edges and discordant edges (those whose incident vertices

[5] $\delta > 0$ is necessary to allow any existing vertex, even one with no in-neighbors, to be chosen as an in-neighbor of a new node.

have different opinions). As described in Section 3.3, a discordant edge (u, v) is either an example of Successful Resistance ($A_u = c, A_v = c'$) or a Failure to Enlighten ($A_u = c', A_v = c$). We declare the discordant edge criterion to be successful if the majority of discordant edges are examples of Successful Resistance. Therefore, in a directed graph, there are three criteria to evaluate. A tie within a criterion causes that criterion to be discarded. If the remaining criteria agree on a result, then that result is declared. Otherwise, if all criteria are discarded, or if the criteria have conflicting outputs, then without further assumptions about the underlying probabilities, the mechanism must output "Undecided".

4.4 Experiment Design

We investigate the performance of the two aggregation methods, across the four random graph models, and a variety of parameters. In each experiment, n_c of the population of n voters are initialized to the correct state c, and then k voters will change their opinions in sequence, as selected by our influence dynamics model.

We set the parameter $\delta = 1$, so that in the directed Barabási-Albert network, each new node has half the chance of connecting to an existing in-neighbor-less node as a node that has in-degree one. We examine our models on graphs of sizes $n \in \{20, 40\}$, with $n_c \in [0, n/2]$ and $k \in [0, 3n/4]$. The simulation was implemented using Python version 3.3.2. Each data point is the result of 1,000 replications at each combination of parameter settings.

5. RESULTS

We measured the performance of our models in terms of the improvement in ability to accurately predict the correct winner. Since our models cannot decide a winner on every possible graph, we compute the performance of the Righteous Argument Model by deciding ties uniformly at random. We compare performance of our model to the performance of the naive model, by defining acc_{naive} to be the accuracy of the naive model, acc_{sg} to be the accuracy of our model, and $\Delta_{sg,naive} = acc_{sg} - acc_{naive}$ to be the *performance improvement* from using our model instead of the naive model. For instance, a value of 0.05 indicates that our model is 5% more accurate (in absolute difference) in deciding the outcome of the election.

The results of our experiments are summarized in Figure 2. The rows of Figure 2 depict sets of heatmaps showing every combination of parameter settings for n_c and k considered. Each map is labeled with the graph generation algorithm used. The horizontal axis of each graph shows increasing values of k (the number of votes flipped by opinion dynamics), while the vertical axis shows increasing values of n_c (the proportion of voters initially holding the correct opinion). The color of the cells corresponds to $\Delta_{sg,naive}$, and darker cells indicate greater advantage for our model. A red X indicates parameter values where the naive model performs better than the new model.

The first row of the results shows performance differences on randomly generated undirected BA graphs, with different graph attachment parameters m, and number of voters n held constant. In all cases, there exists a white region in the top left of the diagram where neither the counts of the votes nor the edges give the correct winner, because all voters begin with the incorrect opinion and have no opportunity to change. In this case, the performance of both models is exactly zero. In the lower right portion of each graph, a similar region exists where, after opinion dynamics, essentially every voter on every run has the correct opinion, so both the naive and Righteous Argument models give the correct answer. In some of the more extreme parameter settings the Naive model actually performs better, but this effect is very small (typically around 1-2%), and not significantly different from zero. In between however, there is a band where the Righteous Argument enjoys a considerable (roughly 20%) advantage. The band is wider when the degree distribution of the graph is relatively flatter.

The second row of results shows similar heatmaps for ER graphs. Here, different values of n are used for the left and center maps (respectively 20 and 40), and exhibit behavior similar to that in the BA graph. The rightmost graph shows performance if counting concordant edges is used as a predictor alone, rather than in combination with the count of the votes as in the Righteous Argument Model.

The third row of results shows heatmaps for directed BA graphs with attachment parameters $m = 2$, $m = 3$, and $m = 4$ from left to right. Here we observe exceptional advantages for the Righteous Argument model in the upper left region, but also some performance deficiencies in the lower right region. Since these performance differences are not present in the undirected case, and the naive model does not make use of edge information, we attribute this to the additional count of Failed Enlightenments and Successful Resistance cases (r and $\dot{r}$) in the model. This is discussed in some detail in the next section.

The fourth row shows maps for directed ER graphs, with broadly similar results to the directed BA case, though with less extreme values.

6. DISCUSSION

Overall our findings confirm the utility of incorporating social network structure into elections. In the undirected case the model's performance improvement on the class of simulated networks arises primarily because of the homophily generated by that process: when many votes are flipped with a bias toward those with many correct neighbors, a disproportionately large number of concordant edges are created for the correct winner. More formally, since the model decides ties randomly, then it will be right on all cases where the models agree, and half of those where they disagree. It is easy to show that this means the model's performance is an average of the performance of just counting the edges and just counting the votes respectively. This is illustrated by the rightmost figure in the undirected ER case (second row of Figure 2), which is simply a linear transform of the Righteous Argument model's performance (middle of second row of Figure 2). We can thus interpret the Righteous Argument model as a "safe" choice. More formally, if the values of $\frac{p}{1-p}$ and $\frac{q}{\dot{q}}$ are unknown, then the model avoids the worst case performance of selecting based on a single criteria. We note further that, using a tie-breaking rule that is sensitive to the specific graph structure considered could also provide significant advantages.

Despite demonstrating conclusive advantages for the undirected case, the new model's benefits on directed graphs appear mixed. This appears to result from a failure to constrain the ratio of $\frac{r}{\dot{r}}$ relative to the ratio of $\frac{q}{\dot{q}}$, which means that the model is vulnerable to making poor decisions on the basis of meager evidence.

In particular, when the opinions of voters are highly homogeneous, the vote and edge counts used by the Righteous Argument model will nearly always agree. As an extreme example, when no voter starts with the correct opinion, and only a single flip is made in the opinion dynamics phase, the correct opinion can not have a majority of the votes, and certainly cannot have a majority of the concordant edges (since there is only one voter with the correct opinion, there are certain to be no correct concordant edges). However, if the flipped voter has more in edges than out edges, the count of successful resistances will be greater than the count of failed conversions for the correct candidate, and not for the incorrect one. This causes the Righteous Argument model to decide using its tie breaking rules instead. This explains the relatively po-

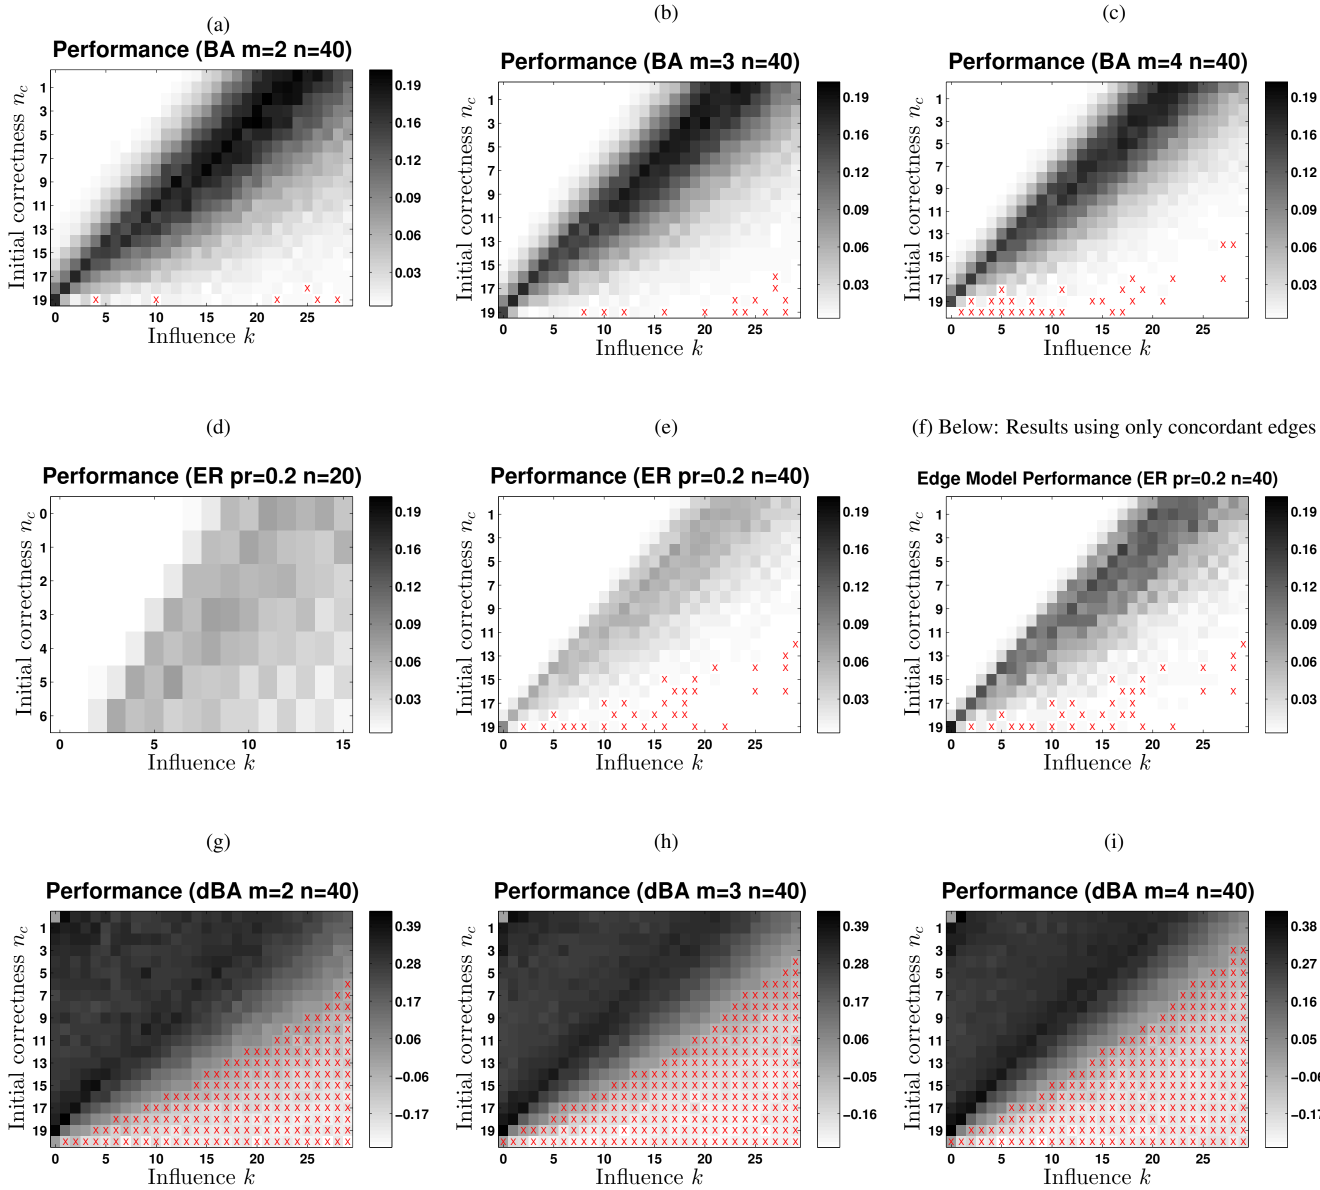

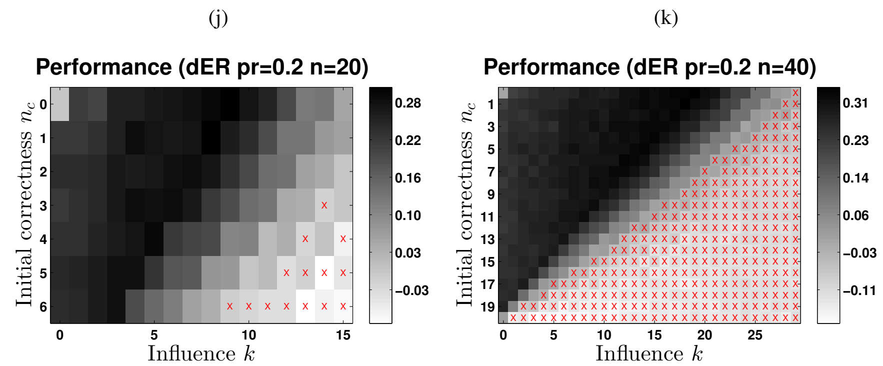

Left and Center: Heatmaps summarizing our results of our voting model. Note that (f) depicts results using only the concordant edge model. Darker colors indicate relative advantage for our model over the naive competitor. The scales to the right of each map show the values corresponding to different colors. Note that (a)-(f) use the same color scale, while the colors in (g)-(k) are only comparable within the same map. Cells with negative values are marked with a red X.

Figure 2

larized performance of the model in the directed case.

The addition of a further constraint on the model ($\frac{r}{\bar{r}} < \frac{q}{\bar{q}}$) allows for a consistent estimator that would avoid these erratic performance issues on some classes of graphs. This estimator would still be efficient to compute, but would under-perform disproportionately on social networks which did not satisfy the assumption.

7. CONCLUSIONS

In this paper, we proposed a new maximum likelihood approach to finding the winner in elections in the presence of a social network. Our new model started from the assumption of righteous argument: that it is easier to convince someone of the truth, than of a falsehood. We derived efficiently computable maximum likelihood estimators for the correct judgment under this assumption for the case of a binary election over an undirected graph, and then extended the model to di-graphs and elections with more than two alternatives. We validated the new approach with extensive simulations over many parameter settings of four different graph generation algorithms, and found significant advantages to incorporating the information present in the structure of the social network, under the righteous argument assumption.

We were also able to characterize the properties of graphs on which our model sees the largest performance advantages, in terms of the expected margin of victory for the different components of the model. The model will outperform naive vote counting when the social network exhibits homophily, but will perform less well when that assumption is violated (which corresponds to the new model's assumption being violated). We also detected situations where the model performed less well on directed graphs, and explained how this behavior could be reconciled with the addition of another constraint into the MLE.

Overall, we have demonstrated that there exists a novel and efficient model for incorporating information about the structure of social networks into social choice problems, which offers increased confidence in the outcomes it predicts over not incorporating the structure of the social network.

8. FUTURE WORK

The proposed model offers several interesting avenues for future work. The Righteous Argument model is a general model, that works over a large family of problems. However, additional performance advantages may be possible over a more restricted family of problems. For instance, a model that incorporates knowledge of the social dynamics used could consider possible starting positions that generated a particular configuration, in much the same way as Conitzer's edge model.

Further analysis under specific opinion dynamics could also allow for the development of more robust theoretical guarantees, and allow for answering interesting questions in the areas of manipulation and control over elections that take place in the presence of a social network. For instance, if the center can control the length of time for social interactions occur over, within some bounds, can they sway the outcome of a given election?

Another possible direction concerns the proper design and evaluation of a system that incorporates the additional constraints on the ratios of $\frac{r}{\bar{r}}$ and $\frac{q}{\bar{q}}$. The fact that decisions are made on the basis of minimal evidence also suggests that a Bayesian approach might be used to decide ties more effectively in the existing model.

9. REFERENCES

[1] Reka Albert and Albert-Laszlo Barabasi. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47–97, 2002.

[2] Marco Bartolozzi, Derek Bruce Leinweber, and Anthony William Thomas. Stochastic opinion formation in scale-free networks. *Physical Review E*, 72(4):046113, 2005.

[3] Duncan Black, Robert Albert Newing, Iain McLean, Alistair McMillan, and Burt L Monroe. *The theory of committees and elections*. Springer, 1958.

[4] Paolo Boldi, Francesco Bonchi, Carlos Castillo, and Sebastiano Vigna. Voting in social networks. In *Proceedings of the 18th ACM conference on Information and knowledge management*, pages 777–786. ACM, 2009.

[5] Béla Bollobás, Christian Borgs, Jennifer Chayes, and Oliver Riordan. Directed scale-free graphs. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 132–139. Society for Industrial and Applied Mathematics, 2003.

[6] James Coleman, Elihu Katz, and Herbert Menzel. The diffusion of an innovation among physicians. *Sociometry*, pages 253–270, 1957.

[7] Vincent Conitzer. Should social network structure be taken into account in elections? *Mathematical Social Sciences*, 64(1):100–102, 2012.

[8] Vincent Conitzer. The maximum likelihood approach to voting on social networks. In *Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton-13)*, pages 1482–1487. Allerton Retreat Center, Monticello, IL, USA, 2013.

[9] Paul Erdős and Alfréd Rényi. On random graphs. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.

[10] Serge Gaspers, Victor Naroditskiy, Nina Narodytska, and Toby Walsh. Possible and necessary winner problem in social polls. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 613–620. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[11] Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3), 2002.

[12] Ulrich Krause. A discrete nonlinear and non-autonomous model of consensus formation. *Communications in difference equations*, pages 227–236, 2000.

[13] André CR Martins. Continuous opinions and discrete actions in opinion dynamics problems. *International Journal of Modern Physics C*, 19(04):617–624, 2008.

[14] Amirali Salehi-Abari and Craig Boutilier. Empathetic social choice on social networks. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 693–700. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[15] David Strang and Sarah A Soule. Diffusion in organizations and social movements: From hybrid corn to poison pills. *Annual review of sociology*, pages 265–290, 1998.

[16] Alan Tsang and Kate Larson. Opinion dynamics of skeptical agents. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 277–284. International Foundation for Autonomous Agents and Multiagent Systems, 2014.