

Revisiting Opinion Dynamics with Varying Susceptibility to Persuasion via Non-Convex Local Search

T-H. Hubert Chan*
The University of Hong Kong

Zhibin Liang
The University of Hong Kong

Mauro Sozio
LTCI, Télécom ParisTech University

ABSTRACT

We revisit the opinion susceptibility problem that was proposed by Abebe et al. [1], in which agents influence one another's opinions through an iterative process. Each agent has some fixed innate opinion. In each step, the opinion of an agent is updated to some convex combination between its innate opinion and the weighted average of its neighbors' opinions in the previous step. The resistance of an agent measures the importance it places on its innate opinion in the above convex combination. Under non-trivial conditions, this iterative process converges to some equilibrium opinion vector. For the unbudgeted variant of the problem, the goal is to select the resistance of each agent (from some given range) such that the sum of the equilibrium opinions is minimized.

Contrary to the claim in the aforementioned KDD 2018 paper, the objective function is in general non-convex. Hence, formulating the problem as a convex program might have potential correctness issues. We instead analyze the structure of the objective function, and show that any local optimum is also a global optimum, which is somehow surprising as the objective function might not be convex. Furthermore, we combine the iterative process and the local search paradigm to design very efficient algorithms that can solve the unbudgeted variant of the problem optimally on large-scale graphs containing millions of nodes.

CCS CONCEPTS

• **Human-centered computing** → **Social network analysis**; • **Mathematics of computing** → **Graph algorithms**.

KEYWORDS

opinion dynamics, susceptibility to persuasion, non-convex local search

ACM Reference Format:

T-H. Hubert Chan, Zhibin Liang, and Mauro Sozio. 2019. Revisiting Opinion Dynamics with Varying Susceptibility to Persuasion via Non-Convex Local Search. In *Proceedings of the 2019 World Wide Web Conference (WWW '19)*, May 13–17, 2019, San Francisco, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3308558.3313509>

*This research was partially supported by a grant from the PROCORE France-Hong Kong Joint Research Scheme sponsored by the Research Grants Council of Hong Kong and the Consulate General of France in Hong Kong under the project F-HKU702/16. T-H. Hubert Chan was partially supported by the Hong Kong RGC under the grant 17200817.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '19, May 13–17, 2019, San Francisco, CA, USA

© 2019 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-6674-8/19/05.

<https://doi.org/10.1145/3308558.3313509>

1 INTRODUCTION

We revisit the opinion susceptibility problem that was proposed by Abebe et al. [1], in which agents in a network influence one another's opinions through an iterative process. In this model, the opinion of an agent at each step is a function of its *innate* opinion and its neighbors' opinions in the previous step. The susceptibility of an agent i to persuasion is measured by a resistance parameter $\alpha_i \in (0, 1)$, where a larger value of α_i means that agent i places more importance on its innate opinion. Under very mild conditions (such as every $\alpha_i > 0$), the iterative process converges to a unique equilibrium opinion vector.

As remarked in [1], opinion interference in social networks have wide applications in product marketing and political campaigns. While most previous works [6, 11, 16] investigated models on changing the opinions directly, [1] is the first work that considered modifying the susceptibility of individual agents. Using a number in $[0, 1]$ to represent an agent's opinion, they considered the problem of modifying the agents' resistances such that the sum of the equilibrium opinions is minimized (or equivalently maximized by applying the transformation $x \mapsto 1 - x$ on each opinion).

Abebe et al. [1] pointed out several works in social psychology that have investigated people's susceptibility to persuasion, including studies on how peer pressure can make adolescents exhibit risky and antisocial behavior [2, 5, 7], and on the role of susceptibility in politics [19]. Factors such as self-esteem, locus of control, awareness of specific issues, as well as exposure to people in the same network can contribute to an individual's susceptibility to persuasion. As a result, by affecting these factors or introducing persuasive cues [15], one can make an individual more or less susceptible to persuasion, i.e., the resistance parameter can be lowered or increased.

Observe that the problem is trivial if the resistance of each agent can be picked from the closed interval $[0, 1]$. For minimizing the equilibrium opinions, it suffices to make the agent with the minimum innate opinion the most resistant (setting its resistance to 1) and everyone else totally compliant (setting its resistance to 0). The problem is non-trivial if the resistance α_i of each agent i can take value from some interval $[l_i, u_i]$, where $0 < l_i < u_i < 1$. Under this restriction, Abebe et al. [1] showed that it suffices to consider extreme points in the resistance vector space, i.e., for each agent i , one can pick $\alpha_i \in \{l_i, u_i\}$. Moreover, they attempt to model the problem using a convex program. However, contrary to their claim, we have discovered that the objective function is in general neither convex nor concave (see Figure 1 in Section 2). Hence, there could be potential correctness issues with their approach. Furthermore, a general convex program is quite expensive to solve, and indeed, the experiments in [1] were performed on networks with less than 1000 nodes, most likely because their approach is non-scalable to large networks.

This work resolves both the correctness and scalability issues in [1]. We give a thorough analysis of the problem structure to design a provably correct local search framework, and further exploit the properties of the iterative process to implement our algorithms that are efficient enough to run on networks with millions of nodes.

Our Contributions. We study the structure of the opinion susceptibility problem [1] in detail, and also make algorithmic contributions, which we describe as follows.

- *Analysis of problem structure.* Contrary to the claim in [1], we discover that the objective function is in general neither convex nor concave. However, we have also recovered (in Corollary 3.4) that for each agent, it is sufficient to set its resistance to either the lower or the upper bound, i.e., we can search for an optimal resistance vector from the set of extreme points of its feasible region.
We analyze the mathematical structure of the problem in Section 3. Perhaps the most important technical insight in this paper is that we show (in Lemma 3.9) that if the current vector solution is not optimal, then there exists a coordinate that can be flipped such that the objective will be strictly improved. This shows that an optimal vector can be found by a simple local search algorithm.
- *Local search with irrevocable updates.* In general, local search could still take exponential time to find an optimal solution, for instance, the simplex algorithm for linear programming. For minimizing the sum of equilibrium opinions, we show (in Lemma 4.1) that starting from the upper bound resistance vector, then the local search algorithm will flip each coordinate at most once, which implies that an optimal vector can be found in polynomial time.
- *Efficient Local Search on Large-Scale Graphs.* Typically, in local search, the objective function needs to be evaluated at the current solution in each step. However, since the objective function involves matrix inverse, its evaluation will be too expensive when the dimension of the matrix is in the order of millions. Instead, we use the iterative process of the opinion dynamics model itself to approximate the equilibrium vector. We have developed several update strategies for local search. For *conservative* or *opportunistic* updates, one always makes sure that the error of the estimated equilibrium vector is small enough before any coordinate of the resistance vector is flipped. For *optimistic* update, one might flip a coordinate of the resistance vector even before the estimated equilibrium vector is accurate enough. However, this might introduce mistakes which need to be corrected later. Nevertheless, experiments show that mistakes are rarely made by the optimistic update strategy. In any case, for all three update strategies, an optimal vector will be returned when the local search terminates.

Our approaches are scalable and can run on networks with millions of nodes. We report the experimental results in Section 5. In particular, using multiple number of threads, the optimistic update strategy can solve the problem optimally on networks with up to around 65 million nodes.

1.1 Related Works

The most related work is the KDD 2018 paper [1] by Abebe et al., who gave a history of the problem and other related works, to which the reader can refer. Here we just give a brief summary.

History of Model. As mentioned in [1], the opinion dynamics model is inspired by the works of DeGroot [4] and Friedkin and Johnsen [8]. DeGroot considered the special case where all resistance α_i 's equal 0, and the opinion vector update rule reduces to $z^{(t)} := Pz^{(t-1)} \in \mathbb{R}^V$, for some row stochastic matrix P . This model is well understood through the theory of Markov chains, which typically uses the transpose of P in the update rule: $z^{(t)} := P^T z^{(t-1)} \in \mathbb{R}^V$. The iterative process considered in this paper is exactly the model in [8], which also investigated the conditions under which the process converges to a unique equilibrium vector.

Budgeted Variant. We have so far discussed only the *unbudgeted* variant of the opinion optimization problem, where the resistance of every agent can be modified. The *budgeted* variant of the problem has also been considered in [1], where some initial resistance vector $\alpha \in (0, 1)^V$ and some budget $k \in \mathbb{Z}$ are given, and the algorithm is allowed to change at most k coordinates of α . It is shown that the budgeted variant is NP-hard [1]. An obvious approach is to use some greedy paradigm to pick k coordinates to modify one by one, where the next coordinate is picked to have the largest improvement in the objective function. However, no theoretical guarantee was given for budgeted variant in [1]. Since our analysis of the unbudgeted variant is already quite involved, this paper will focus on the unbudgeted case. We will leave the analysis of the greedy algorithm for the budgeted variant as future work, which we briefly discuss in the conclusion section.

Non-Convex Optimization. In general, optimizing a non-convex function under non-convex constraints is NP-hard. However, in many cases, one can exploit the structure of the objective function or constraints to devise polynomial-time algorithms; see the survey by Jain and Kar [13] on non-convex optimization algorithms encountered in machine learning. Indeed, variants of the gradient descent have been investigated to escape saddle points by Jin et al. [14], who also gave examples of problems where all local optima are also global optima; some examples are tensor decomposition [9], dictionary learning [22], phase retrieval [21], matrix sensing [3, 20] and matrix completion [10]. However, all these problems involve some quadratic loss functions, whose structures are totally different from our objective functions which involve matrix inverse.

Hartman [12] considered the special case that the objective function is the difference of two convex functions. Strelakovsky devised a local search method to optimize such objective functions. Even though the objective functions in our problem are somewhere convex and somewhere concave (see Figure 1), it is not immediately clear if they can be expressed as differences of convex functions.

2 PRELIMINARIES

We revisit the opinion susceptibility problem that was proposed by Abebe et al. [1], which we rephrase using more general parameters. In an opinion dynamics model, there is a set V of agents, whose innate opinions are given by a vector $s \in [0, 1]^V$. The interaction between agents are captured by a row stochastic matrix $P \in [0, 1]^{V \times V}$, i.e., each entry of P is non-negative and every row sums to 1. In

the rest of this paper, we always call P the interaction matrix. The susceptibility of the agents is measured by the resistance vector $\alpha \in (0, 1)^V$, where a higher α_i value means that agent i is more resistant.

As mentioned in [1], the opinion vector of agents evolves in discrete time according to the equation $z^{(t+1)} := As + (I - A)Pz^{(t)}$, where $A = \text{Diag}(\alpha)$ is the diagonal matrix with $A_{ii} = \alpha_i$ and I is the identity matrix. Equating $z^{(t)}$ with $z^{(t+1)}$, one can see that the equilibrium opinion vector is given by $z = [I - (I - A)P]^{-1}As$, which exists under non-trivial conditions such as every $\alpha_i > 0$. The objective is to choose a resistance vector α to minimize the sum of equilibrium opinions $\langle \mathbf{1}, z \rangle = \mathbf{1}^T z$. Observe that one can also consider maximizing the sum of equilibrium opinions; however, since the techniques are essentially the same, we will focus on the minimization variant of the problem.

Definition 2.1 (Opinion Susceptibility Problem). Given a set V of agents with innate opinions $s \in [0, 1]^V$ and interaction matrix $P \in [0, 1]^{V \times V}$, suppose for each $i \in V$, its resistance is restricted to some interval $I_i := [l_i, u_i] \subseteq [0, 1]$ where we assume that $0 < l_i < u_i < 1$.

The objective is to choose $\alpha \in I_V := \times_{i \in V} I_i \subseteq [0, 1]^V$ such that the following objective function is minimized:

$$f(\alpha) := \mathbf{1}^T [I - (I - A)P]^{-1}As,$$

where $A = \text{Diag}(\alpha)$ is the diagonal matrix with $A_{ii} = \alpha_i$. Observe that the assumption $\alpha > 0$ ensures that the above inverse exists.

Unbudgeted vs Budgeted Variants. In Definition 2.1, we are allowed to modify the resistance of any agent, and this is known as the unbudgeted variant. However, in [1], the budgeted variant is also considered: given some initial resistance vector and a budget k , the resistance of at most k agents can be changed. In this paper, we focus on efficient algorithms that optimally solve the unbudgeted variant.

Technical Assumption. To simplify our proofs, we assume that the interaction matrix P corresponds to an irreducible random walk. Irreducibility is satisfied if P arises from a connected graph.

Non-convex Objective. Contrary to the claim in [1], the objective f in Definition 2.1 is in general not a convex function of α . In fact, the following example shows that it might be neither convex nor concave. Consider three vertices $V = \{1, 2, 3\}$, where the innate

vector s and the interaction matrix P are given by: $s = \begin{bmatrix} 1 \\ 0.5 \\ 0 \end{bmatrix}$ and

$$P = \begin{bmatrix} 0 & 0.5 & 0.5 \\ 0.5 & 0 & 0.5 \\ 0.5 & 0.5 & 0 \end{bmatrix}.$$

Suppose we fix $\alpha_2 = \alpha_3 = 0.1$ and consider the objective as a function of α_1 as $g(\alpha_1) = \mathbf{1}^T [I - (I - A)P]^{-1}As$, where $A = \text{Diag}(\alpha_1, \alpha_2, \alpha_3)$. Then, the plot of g in Figure 1 (a) shows that it is not convex. Moreover, suppose this time we fix $\alpha_1 = \alpha_2 = 0.1$ and consider the objective as a function of α_3 as $h(\alpha_3) = \mathbf{1}^T [I - (I - A)P]^{-1}As$. Then, the plot of h in Figure 1 (b) shows that it is not concave.




Figure 1: Cross-Sections of Objective Function

3 STRUCTURAL PROPERTIES OF OBJECTIVE FUNCTION

In this section, we investigate the properties of the objective function f in Definition 2.1; we assume that the interaction matrix P and the innate opinion vector s are fixed, and f is a function on the resistance vector α . Contrary to the claim in [1], the objective function is in general not convex. Fortunately, we can still exploit some properties of the function. As we shall see, even when the function is not convex, every local optimum (which will be defined formally) is a global optimum. This enables us to use variants of the local search method to solve the problem optimally.

3.1 Marginal Monotonicity

As in [1], we show that when one chooses the resistance α_i for each agent $i \in V$, it suffices to consider the extreme points $\{l_i, u_i\}$. Our approach explicitly analyzes the partial derivative $\frac{\partial f(\alpha)}{\partial \alpha_i}$ which plays a major role in the local search algorithm that we later develop.

Intuition: Guidance by Current Equilibrium Vector. Observe that given the innate opinion vector s and irreducible interaction matrix P , for some resistance vector $\alpha \in (0, 1)^V$, the equilibrium opinion vector is given by $z(\alpha) = [I - (I - A)P]^{-1}As$, where $A = \text{Diag}(\alpha)$. For some $i \in V$, if the innate opinion s_i is larger than its equilibrium $z_i(\alpha)$, this suggests that by being more stubborn, agent i should be able to increase its equilibrium opinion. In other words, one would expect $\frac{\partial z_i(\alpha)}{\partial \alpha_i}$ and $s_i - z_i(\alpha)$ to have the same sign. However, what is surprising is that in Lemma 3.2, we shall see that even for any $j \in V$, $\frac{\partial z_j(\alpha)}{\partial \alpha_i}$ and $s_i - z_i(\alpha)$ have the same sign.

Notation. For any $\alpha \in \mathbb{R}^V$ and $K \subseteq V$, let $\alpha_{-K} \in \mathbb{R}^V$ denote the vector such that $\alpha_{-K}(i) = \alpha(i)$ if $i \notin K$, and $\alpha_{-K}(i) = 0$ if $i \in K$, i.e., the coordinates K of α are replaced with 0. Similarly, given $\alpha \in \mathbb{R}^V$, we denote $A_{-K} = \text{Diag}(\alpha_{-K})$.

In Definition 2.1, observe that the inverse $[I - (I - A)P]^{-1}$ is involved in the objective function $f(\alpha)$, where $A = \text{Diag}(\alpha)$. Since we wish to analyze the effect on $f(\alpha)$ of changing only a subset of coordinates in α , the next lemma will be used for simplifying matrix arithmetic involving the computation of inverses. Its proof is deferred to Section 6.

LEMMA 3.1 (INVERSE ARITHMETIC). Given $K \subseteq V$ and $\alpha \in (0, 1)^V$, let $A := \text{Diag}(\alpha)$ and recall that P is the irreducible interaction matrix. Then, the inverse $M = [I - (I - A_{-K})P]^{-1}$ exists, and every entry

of M is positive. Moreover, for each $k \in V$, define $a_k = 0$ if $k \in K$, otherwise $a_k = \alpha_k$. Then, we have:

- (1) $(PM)_{kk} = \frac{M_{kk}-1}{1-a_k} > 0$;
- (2) $(PM)_{kj} = \frac{M_{kj}}{1-a_k} > 0$, for each $j \neq k$.

LEMMA 3.2 (SIGN OF PARTIAL DERIVATIVE). *In the Opinion Susceptibility Problem in Definition 2.1, given the innate opinion vector s and irreducible interaction matrix P , recall that $z(\alpha) := [I - (I - A)P]^{-1}As$, where $A = \text{Diag}(\alpha)$. Then, for any $\alpha \in (0, 1)^V$ and any $i, k \in V$, the two values $\frac{\partial z_k(\alpha)}{\partial \alpha_i}$ and $s_i - z_i(\alpha)$ have exactly the same sign in $\{-, 0, +\}$.*

In particular, this implies that $\frac{\partial f(\alpha)}{\partial \alpha_i} = \sum_{k \in V} \frac{\partial z_k(\alpha)}{\partial \alpha_i}$ also has the same sign as $s_i - z_i(\alpha)$.

PROOF. By the definition of the inverse of a matrix B , we have $BB^{-1} = I$. The partial derivative with respect to a variable t is: $\frac{\partial B}{\partial t}B^{-1} + B\frac{\partial B^{-1}}{\partial t} = 0$. Hence, we have $\frac{\partial B^{-1}}{\partial t} = -B^{-1}\frac{\partial B}{\partial t}B^{-1}$. Applying the above result with $B = I - (I - A)P$ and $t = \alpha_i$ and denoting $M = [I - (I - A)P]^{-1}$, we get $\frac{\partial M}{\partial \alpha_i} = -Me_i e_i^T PM$. Considering $z(\alpha) = MAs$, we have $\frac{\partial z(\alpha)}{\partial \alpha_i} = \frac{\partial M}{\partial \alpha_i}As + Me_i e_i^T s$. Replacing $\frac{\partial M}{\partial \alpha_i}$, we obtain for any $i, k \in V$:

$$\frac{\partial z_k(\alpha)}{\partial \alpha_i} = -e_k^T Me_i e_i^T PMAs + e_k^T Me_i e_i^T s = M_{ki} \cdot [s_i - e_i^T Pz(\alpha)].$$

By Lemma 3.1 with $K = \emptyset$, we know that every entry of M is positive. Thus, the sign of $\frac{\partial z_k(\alpha)}{\partial \alpha_i}$ is the same as that of the scalar $s_i - e_i^T Pz(\alpha)$.

Recalling $M = [I - (I - A)P]^{-1}$, we have $[I - (I - A)P]M = I \Rightarrow (I - A)PM = M - I \Rightarrow PM = (I - A)^{-1}(M - I)$ where $(I - A)^{-1}$ exists since $\alpha_j < 1$ for each $j \in V$.

Next, since $z(\alpha) = MAs$, we have:

$$Pz(\alpha) = PMAs = (I - A)^{-1}(M - I)As = (I - A)^{-1}[z(\alpha) - As].$$

Finally, replacing $Pz(\alpha)$, we have

$$\begin{aligned} s_i - e_i^T Pz(\alpha) &= s_i - e_i^T (I - A)^{-1}[z(\alpha) - As] \\ &= s_i - \frac{1}{1 - \alpha_i} [z_i(\alpha) - \alpha_i s_i] \\ &= \frac{1}{1 - \alpha_i} [s_i - z_i(\alpha)]. \end{aligned}$$

Since $1 - \alpha_i > 0$, it follows that $\frac{\partial z_k(\alpha)}{\partial \alpha_i}$ and $s_i - z_i(\alpha)$ have exactly the same sign in $\{-, 0, +\}$, as required. \square

The next lemma shows that the sign of the partial derivatives with respect to coordinate i is actually independent of the current value α_i . Its proof is deferred to Section 6.

LEMMA 3.3 (SIGN OF PARTIAL DERIVATIVE INDEPENDENT OF COORDINATE VALUE). *Referring to Lemma 3.2. For any $\alpha \in (0, 1)^V$ and any $i \in V$, denote $M = [I - (I - A_{-\{i\}})P]^{-1}$. Then, $\frac{\partial f(\alpha)}{\partial \alpha_i}$ has the same sign in $\{-, 0, +\}$ as $s_i - \sum_{j \neq i} M_{ij}\alpha_j s_j$, which is independent of α_i .*

COROLLARY 3.4 (EXTREME POINTS ARE SUFFICIENT). *In Definition 2.1, for any $i \in V$, fixing the resistance values of all other agents except i , the objective $f(\alpha)$ is a monotone function in α_i . This implies that*

to minimize f , it suffices to consider the extreme points $\alpha_i \in \{l_i, u_i\}$, for each $i \in V$.

3.2 Local vs Global Optimum

As shown in Corollary 3.4, it suffices to choose the resistance vector α from the extreme points in Definition 2.1. Lemma 3.2 readily gives a method to decide, given a current choice of α , whether the objective f can be decreased by changing the resistance of some agent. In Lemma 3.9, we show that if α is not a global minimum, then such an agent must exist. As we shall see, this implies that a local search method can find a global minimum.

Given α and $\alpha' \in \mathbb{R}^V$, denote $\Delta(\alpha, \alpha') := \{i \in V : \alpha_i \neq \alpha'_i\}$ as the set of coordinates at which the vectors differ.

Definition 3.5 (Local Minimizer). Given an objective function $f : \mathcal{I}_V \rightarrow \mathbb{R}$, a vector $\alpha \in \mathcal{I}_V$ is a local minimizer of f , if for all $\alpha' \in \mathcal{I}_V$ such that $\Delta(\alpha, \alpha') = 1$, $f(\alpha) \leq f(\alpha')$.

Notation. When we wish to consider the effect of changing the resistance of only 2 agents $i \neq k \in V$, we write $f(\alpha) = f(\alpha_i, \alpha_k)$, assuming that $\alpha_{-\{i,k\}}$ is fixed.

Lemmas 3.6 and 3.7 give some technical results involving changing the resistance of two agents. Their proofs are deferred to Section 6.

LEMMA 3.6. *For any $i, k \in V$ such that $i \neq k$, let $M = [I - (I - A_{-\{i\}})P]^{-1}$ and $R = [I - (I - A_{-\{i,k\}})P]^{-1}$. Then for any $j \in V$, we have*

- (1) $M_{jk} = \frac{R_{jk}}{1 + \alpha_k R_{kk} - \alpha_k}$,
- (2) $M_{jh} = R_{jh} - \frac{\alpha_k R_{jk} R_{kh}}{1 + \alpha_k R_{kk} - \alpha_k}$, for $h \neq k$.

In particular, the quantity in Lemma 3.3 can be rewritten as follows:

$$s_i - \sum_{j \neq i} M_{ij}\alpha_j s_j = s_i - \sum_{j \neq i, k} R_{ij}\alpha_j s_j - \frac{\alpha_k R_{ik}}{1 + \alpha_k R_{kk} - \alpha_k} (s_k - \sum_{j \neq i, k} R_{kj}\alpha_j s_j).$$

LEMMA 3.7 (DIAGONAL ENTRY). *Suppose $\alpha \in (0, 1)^V$, recall that $A_{-\{i,k\}} := \text{Diag}(\alpha_{-\{i,k\}})$, and P corresponds to an irreducible interaction matrix. For any $i, k \in V$ such that $i \neq k$, let $R = [I - (I - A_{-\{i,k\}})P]^{-1}$, then $R_{ii} = \max_{j \in V} R_{ji}$. Moreover, $R_{ii} = R_{ki}$ if and only if $P_{kk} + P_{ki} = 1$.*

The following lemma gives the key insight for why local search works. Intuitively, it shows that there does not exist any discrete "saddle point". Even though its proof is technical, we still include it here because of its importance.

LEMMA 3.8 (SWITCHING LEMMA). *Recall that f is defined in Definition 2.1 with an irreducible interaction matrix P , and assume $|V| \geq 3$. Suppose $\alpha, \beta \in (0, 1)^V$ such that $\Delta(\alpha, \beta) = \{i, k\}$ for some $i \neq k$. Moreover, suppose further that*

$$\min\{f(\alpha_i, \alpha_k), f(\beta_i, \beta_k)\} < \min\{f(\alpha_i, \beta_k), f(\beta_i, \alpha_k)\}.$$

Then, we have

$$\max\{f(\alpha_i, \alpha_k), f(\beta_i, \beta_k)\} > \min\{f(\alpha_i, \beta_k), f(\beta_i, \alpha_k)\}.$$

PROOF. We prove the lemma by contradiction. Suppose

$$\max\{f(\alpha_i, \alpha_k), f(\beta_i, \beta_k)\} \leq \min\{f(\alpha_i, \beta_k), f(\beta_i, \alpha_k)\}.$$

Without loss of generality, suppose further that $f(\alpha_i, \alpha_k) \geq f(\beta_i, \beta_k)$. Then, we have

$$f(\alpha_i, \alpha_k) \leq \begin{cases} f(\beta_i, \alpha_k) \\ f(\alpha_i, \beta_k) \end{cases} \quad \text{and} \quad f(\beta_i, \beta_k) < \begin{cases} f(\beta_i, \alpha_k) \\ f(\alpha_i, \beta_k) \end{cases}.$$

We remark that it is important to distinguish between the strict and non-strict inequality. We use the notation f'_i to denote the partial derivative with respect to coordinate i .

From $f(\alpha_i, \alpha_k) \leq f(\beta_i, \alpha_k)$ and the fact that f is marginally monotone (Lemma 3.3) and $f'_i(x, \alpha_k)$ has the same sign in $\{-, 0, +\}$ for $x \in (0, 1)$, we have

$$f'_i(_, \alpha_k) \cdot (\alpha_i - \beta_i) \leq 0. \quad (1)$$

On the other hand, from the strict inequality $f(\beta_i, \beta_k) < f(\beta_i, \alpha_k)$, we know the partial derivative $f'_k(\beta_i, y)$ must have the same non-zero sign in $\{-, +\}$, again from Lemma 3.3. Therefore, we have:

$$f'_k(\beta_i, _) \cdot (\alpha_k - \beta_k) > 0. \quad (2)$$

Similarly, $f(\alpha_i, \alpha_k) \leq f(\alpha_i, \beta_k)$ and $f(\beta_i, \beta_k) < f(\alpha_i, \beta_k)$ give the following:

$$f'_k(\alpha_i, _) \cdot (\alpha_k - \beta_k) \leq 0, \quad (3)$$

$$f'_i(_, \beta_k) \cdot (\alpha_i - \beta_i) > 0. \quad (4)$$

Next, using Lemma 3.3 and R as defined in Lemma 3.6, the above inequalities (1) to (4) become:

$$[s_i - \sum_{j \neq i, k} R_{ij} \alpha_j s_j - \frac{\alpha_k R_{ik}}{1 + \alpha_k R_{kk} - \alpha_k} (s_k - \sum_{j \neq i, k} R_{kj} \alpha_j s_j)] (\alpha_i - \beta_i) \leq 0,$$

$$[s_k - \sum_{j \neq i, k} R_{kj} \beta_j s_j - \frac{\beta_i R_{ki}}{1 + \beta_i R_{ii} - \beta_i} (s_i - \sum_{j \neq i, k} R_{ij} \beta_j s_j)] (\alpha_k - \beta_k) > 0,$$

$$[s_k - \sum_{j \neq i, k} R_{kj} \alpha_j s_j - \frac{\alpha_i R_{ki}}{1 + \alpha_i R_{ii} - \alpha_i} (s_i - \sum_{j \neq i, k} R_{ij} \alpha_j s_j)] (\alpha_k - \beta_k) \leq 0,$$

$$[s_i - \sum_{j \neq i, k} R_{ij} \beta_j s_j - \frac{\beta_k R_{ik}}{1 + \beta_k R_{kk} - \beta_k} (s_k - \sum_{j \neq i, k} R_{kj} \beta_j s_j)] (\alpha_i - \beta_i) > 0.$$

Recall that $\alpha_j = \beta_j$ for $j \neq i, k$. Hence, we denote:

$$c_i := s_i - \sum_{j \neq i, k} R_{ij} \alpha_j s_j = s_i - \sum_{j \neq i, k} R_{ij} \beta_j s_j,$$

$$c_k := s_k - \sum_{j \neq i, k} R_{kj} \alpha_j s_j = s_k - \sum_{j \neq i, k} R_{kj} \beta_j s_j,$$

$$g_i(x) := \frac{x R_{ki}}{1 + x R_{ii} - x} \quad \text{and} \quad g_k(x) := \frac{x R_{ik}}{1 + x R_{kk} - x}.$$

Then, we have

$$[c_i - g_k(\alpha_k) c_k] (\alpha_i - \beta_i) \leq 0, \quad (5)$$

$$[c_k - g_i(\beta_i) c_i] (\alpha_k - \beta_k) > 0, \quad (6)$$

$$[c_k - g_i(\alpha_i) c_i] (\alpha_k - \beta_k) \leq 0, \quad (7)$$

$$[c_i - g_k(\beta_k) c_k] (\alpha_i - \beta_i) > 0. \quad (8)$$

Observe that $c_i \neq 0$, otherwise (6) and (7) contradict each other. Similarly, $c_k \neq 0$, otherwise (5) and (8) contradict each other. We next argue that $c_i c_k > 0$.

From (5) and (8) we have

$$[c_i - g_k(\alpha_k) c_k] [c_i - g_k(\beta_k) c_k] \leq 0. \quad (9)$$

If $c_i c_k < 0$, then the above expression will be positive, because $g_k(\cdot) \geq 0$ (we shall see that later). Hence, we conclude that c_i and c_k have the same sign.

From (6) and (7) we have

$$[c_k - g_i(\beta_i) c_i] [c_k - g_i(\alpha_i) c_i] \leq 0. \quad (10)$$

Rearranging (9) and (10), we have:

$$[\frac{c_i}{c_k} - g_k(\alpha_k)] [\frac{c_i}{c_k} - g_k(\beta_k)] \leq 0,$$

and

$$[\frac{c_k}{c_i} - g_i(\beta_i)] [\frac{c_k}{c_i} - g_i(\alpha_i)] \leq 0.$$

Note that every entry of R is positive by Lemma 3.1 and we can easily prove $g_i(\cdot)$ and $g_k(\cdot)$ are both strictly increasing functions in $[0, 1]$. Since $\alpha, \beta \in (0, 1)^V$, the above two inequalities imply that

$$0 = g_k(0) < g_k(\min\{\alpha_k, \beta_k\}) \leq \frac{c_i}{c_k},$$

$$\frac{c_i}{c_k} \leq g_k(\max\{\alpha_k, \beta_k\}) < g_k(1) = \frac{R_{ik}}{R_{kk}} \leq 1,$$

and

$$0 = g_i(0) < g_i(\min\{\alpha_i, \beta_i\}) \leq \frac{c_k}{c_i},$$

$$\frac{c_k}{c_i} \leq g_i(\max\{\alpha_i, \beta_i\}) < g_i(1) = \frac{R_{ki}}{R_{ii}} \leq 1,$$

where $\frac{R_{ik}}{R_{kk}} \leq 1$ and $\frac{R_{ki}}{R_{ii}} \leq 1$ are from Lemma 3.7.

Notice that we get $0 < \frac{c_i}{c_k} < 1$ and $0 < \frac{c_k}{c_i} < 1$, which is a contradiction. Hence, the proof is completed. \square

LEMMA 3.9 (DESCENDING COORDINATE). *Let f be the function as defined in Definition 2.1. Suppose $\alpha, \beta \in (0, 1)^V$ such that $f(\alpha) > f(\beta)$. Then, there exists some $i \in \Delta(\alpha, \beta)$ and $\gamma \in (0, 1)^V$ such that $\Delta(\alpha, \gamma) = \{i\}$, $\Delta(\gamma, \beta) = \Delta(\alpha, \beta) \setminus \{i\}$, and $f(\alpha) > f(\gamma)$.*

In other words, by switching one coordinate (corresponding to i) of α to that of β , the objective function f decreases strictly.

PROOF. We prove the lemma by induction on $|\Delta(\alpha, \beta)|$. The base case $|\Delta(\alpha, \beta)| = 1$ is trivial.

We consider the inductive step with $|\Delta(\alpha, \beta)| = q$, for some $q \geq 2$. Given a list S of coordinates from $\Delta(\alpha, \beta)$, we use $\alpha_{[S]}$ to denote the resulting vector obtained from switching coordinates S of α to those of β .

For contradiction's sake, we assume that for all $j \in \Delta(\alpha, \beta)$, $f(\alpha_{[j]}) \geq f(\alpha)$; moreover, we pick $i \in \Delta(\alpha, \beta)$ such that $f(\alpha_{[i]})$ is minimized.

Observe that $f(\alpha_{[i]}) \geq f(\alpha) > f(\beta)$ and $|\Delta(\alpha_{[i]}, \beta)| = q - 1$. Therefore, by the induction hypothesis, there exists some $k \in \Delta(\alpha_{[i]}, \beta)$ such that $f(\alpha_{[i]}) > f(\alpha_{[i, k]})$.

Next, starting from α , we shall fix all coordinates in V except i and k , and we write the objective $f(x, y)$ as a function on only these two coordinates.

Observe that we have already assumed that

$$f(\alpha_i, \alpha_k) \leq \min\{f(\beta_i, \alpha_k), f(\alpha_i, \beta_k)\}. \quad (11)$$

Moreover, from above, we have $f(\alpha_{[i]}) > f(\alpha_{[i, k]})$, which translates to $f(\beta_i, \alpha_k) > f(\beta_i, \beta_k)$. Observe that we must have $f(\alpha_i, \beta_k) \leq f(\beta_i, \beta_k)$; otherwise, we have $f(\beta_i, \beta_k) < \min\{f(\beta_i, \alpha_k), f(\alpha_i, \beta_k)\}$, which, together with (11), will contradict Lemma 3.8.

Therefore, we have $f(\alpha_{[k]}) = f(\alpha_i, \beta_k) \leq f(\beta_i, \beta_k) < f(\beta_i, \alpha_k) = f(\alpha_{[i]})$, which contradicts the choice of $i \in \Delta(\alpha, \beta)$ to minimize $f(\alpha_{[i]})$. This completes the inductive step and also the proof of the lemma. \square

COROLLARY 3.10. *For the function f in Definition 2.1 every local minimizer is a global minimizer.*

PROOF. Suppose that α is a local minimizer, but there is some β with $f(\alpha) > f(\beta)$. Then, Lemma 3.9 implies there is some γ with $|\Delta(\alpha, \gamma)| = 1$ such that $f(\alpha) > f(\gamma)$, contradicting that α is a local minimizer. \square

4 EFFICIENT LOCAL SEARCH

In Section 3, we conclude in Corollary 3.4 that it suffices to consider the extreme points of the search space of resistance vectors. Moreover, Corollary 3.10 states that every local minimizer is a global minimizer. Since we know how to compute the sign of the partial derivative with respect to each coordinate using Lemma 3.2, we can design a simple local search algorithm to find a global minimizer.

However, it is possible that $O(2^n)$ extreme points are encountered before a global minimizer is reached. Fortunately, in this section, we will explore further properties of the objective function, and design a local search algorithm that encounters at most $O(n)$ extreme points before finding a global minimizer.

4.1 Irrevocable Updates

Local Search Strategy. We shall start with the upper bound resistance vector, i.e., for each $i \in V$, $\alpha_i = u_i$. This induces the corresponding equilibrium opinion vector $z(\alpha)$. According to Lemma 3.2, if there is some agent i such that $\alpha_i = u_i$ and $s_i - z_i(\alpha) > 0$, then we should flip α_i to the lower bound l_i . The following lemma shows that each α_i will be flipped at most once. Essentially, we show that we will never encounter the situation that there is some agent k such that $\alpha_k = l_k$ and $s_k - z_k(\alpha) < 0$, in which case we would have to switch α_k back to u_k .

LEMMA 4.1 (EACH COORDINATE FLIPPED AT MOST ONCE). *Starting from the upper bound resistance vector, the above local search strategy flips each α_i at most once.*

PROOF. We first show that for each agent $k \in V$, the quantity $s_k - z_k(\alpha)$ cannot decrease when α is modified according to the local search strategy. According to the strategy, α is modified because there is some agent i such that $\alpha_i = u_i$ and $s_i - z_i(\alpha) > 0$. By Lemma 3.2, $\frac{\partial z_k}{\partial \alpha_i} > 0$ for each $k \in V$. Hence, after α_i is switched from u_i to l_i , $z_k(\alpha)$ decreases, and the quantity $s_k - z_k(\alpha)$ increases.

Observe that if a coordinate α_k is ever flipped from u_k to l_k , this means that at that moment, we must have $s_k - z_k(\alpha) > 0$, which, as we have just shown, will stay positive after α is subsequently updated. \square

4.2 Approximating the Equilibrium Vector

Observe that in our local search algorithm, we need to compute the equilibrium opinion vector $z(\alpha) = [I - (I - A)P]^{-1}As$ for the current resistance vector α , where $A = \text{Diag}(\alpha)$. However, computing matrix inverse is an expensive operation. Instead, we approximate $z(\alpha)$ using the recurrence $z^{(0)} \in [0, 1]^V$ and $z^{(t+1)} :=$

$As + (I - A)Pz^{(t)}$. The following lemma gives an upper bound on the additive error for each coordinate.

LEMMA 4.2 (APPROXIMATION ERROR). *Suppose for some $\epsilon > 0$, for all $i \in V$, $\alpha_i \geq \epsilon$. Then, for every $t \geq 0$, $\|z(\alpha) - z^{(t)}\|_\infty \leq \frac{(1-\epsilon)^t}{\epsilon}$.*

PROOF. Using the Neumann series $[I - (I - A)P]^{-1} = \sum_{j=0}^{\infty} [(I - A)P]^j$, we have $z(\alpha) - z^{(t)} = \sum_{j=t}^{\infty} [(I - A)P]^j As - [(I - A)P]^t z^{(0)}$.

We next prove, by induction, that for any $x \in [0, 1]^V$, $\|[(I - A)P]^j x\|_\infty \leq (1 - \epsilon)^j$, for all $j \geq 0$. The base case $j = 0$ is trivial because every coordinate of x is between 0 and 1. For the inductive step, assume that for some $j \geq 0$, every coordinate of $y = [(I - A)P]^j x$ has magnitude at most $(1 - \epsilon)^j$. Since P is a row stochastic matrix, it follows that $\|Py\|_\infty \leq (1 - \epsilon)^j$; finally, since $\alpha_i \geq \epsilon$ for all $i \in V$, we have $\|(I - A)Py\|_\infty \leq (1 - \epsilon)^{j+1}$, completing the induction proof.

Finally, observing that both $\sum_{j=t}^{\infty} [(I - A)P]^j As$ and $[(I - A)P]^t z^{(0)}$ have non-negative coordinates, we have

$$\|z(\alpha) - z^{(t)}\|_\infty \leq \max\left\{\left\|\sum_{j=t}^{\infty} [(I - A)P]^j As\right\|, \left\|[(I - A)P]^t z^{(0)}\right\|\right\} \leq \sum_{j=t}^{\infty} (1 - \epsilon)^j = \frac{(1 - \epsilon)^t}{\epsilon}$$

as required. \square

4.3 Local Search Algorithm

Based on Lemmas 4.1 and 4.2, we give a local search framework in Algorithm 1. Observe that in line 1, we perturb the innate opinions s slightly to ensure that for each resistance vector α encountered, no coordinate of s and $z(\alpha)$ would coincide.

The **while** loop in line 4 combines local search to update α and estimation of the equilibrium vector $z(\alpha)$. Here are two general update strategies, which are both captured by the non-deterministic step in line 7:

- **Conservative Update.** The opinion vector z is iteratively updated in line 5 until all coordinates of z and s are sufficiently far apart. Then, for every coordinate α_i such that $\alpha_i = u_i$ and $z_i < s_i$, we flip α_i to the lower bound l_i . After we update α , we reset t to 0, and continue to iteratively update z . Whenever we update α and set t to 0, we say that a new *phase* begins; we use the convention that the initial phase is known as phase 0.
- **Opportunistic Update.** Instead of waiting for the approximation error of every coordinate of z to be small enough, we can update some coordinates α_i if $\alpha_i = u_i$ and $z_i \leq s_i - \text{err}(t)$ is small enough. However, there is some tradeoff between waiting for the errors of all coordinates to be small enough and updating coordinates of α that are ready sooner. In Section 5, we will evaluate empirically different update strategies.

Optimistic Update. In both conservative and opportunistic updates, a coordinate α_i is flipped only when we know for sure that the current estimate z_i has small enough error with respect to the equilibrium $z_i(\alpha)$; hence, no mistake in flipping any α_i is ever made.

Algorithm 1: Local Search Framework

Input: Innate opinions $s \in [0, 1]^V$; interaction matrix P ; for each agent $i \in V$, upper u_i and lower l_i bounds for resistance.
Output: Optimal resistance vector $\alpha \in \times_{i \in V} \{l_i, u_i\}$.

- 1 (Technical step.) Randomly perturb each coordinate of s slightly.
- 2 Initially, for each agent i , set $\alpha_i \leftarrow u_i$ to its upper bound; denote $\epsilon_\alpha := \min_{i \in V} \alpha_i$.
- 3 Pick arbitrary $z \in [0, 1]^V$, and set $t \leftarrow 0$; denote $\text{err}(t) := \frac{(1-\epsilon_\alpha)^t}{\epsilon_\alpha}$.
- 4 **while** $\exists i \in V : |s_i - z_i| \leq \text{err}(t) \vee (z_i < s_i \wedge \alpha_i = u_i)$ **do**
 - 5 $z \leftarrow As + (I - A)Pz$, where $A = \text{Diag}(\alpha)$
 - 6 $t \leftarrow t + 1$
 - 7 (Non-deterministic step.) Pick arbitrary $L \subseteq V$ (that can be empty) such that for each $i \in L$, $z_i \leq s_i - \text{err}(t)$ and $\alpha_i = u_i$.
 - 8 **if** $L \neq \emptyset$ **then**
 - 9 **for each** $i \in L$ **do**
 - 10 $\alpha_i \leftarrow l_i$ to its lower bound (and update ϵ_α).
 - 11 $t \leftarrow 0$
- 12 **return** Resistance vector α .

However, our insight is that as the algorithm proceeds, the general trend is for every z_i to decrease.

The first intuition is that if we start with some $z^{(0)}$ such that every coordinate of $z^{(0)}$ is at least its equilibrium coordinate of $z(\alpha)$, then $z^{(t)}$ should converge to $z(\alpha)$ from above. The second observation is that every time we flip some α_i , this will not increase any coordinate of the equilibrium vector $z(\alpha)$, thereby preserving the condition that the current estimate $z^{(t)} \geq z(\alpha)$. Hence, without worrying about the accuracy of the current estimate z , we will simply flip coordinate α_i to l_i when z_i drops below s_i . However, it is possible that we might need to flip α_i back to u_i , if z_i increases in the next iteration and becomes larger than s_i again. We shall see in Section 5 that this scenario is extremely rare. Specifically, in line 8 of Algorithm 2, the set J is (almost) always empty.

5 EXPERIMENTS

Experimental setup. Our experiments run on a server with 2.1 GHz Intel Xeon Gold 6152 CPU and 64GB of main memory. The server is limited to activate at most 24 threads by the administrator. The real network topologies we use in our experiment are shown in Table 1; we interpret each network as an undirected graph. The number n of nodes in the dataset networks ranges from about 1 million to 65 million; in each network, the number m of edges is around $2n$ to $30n$.

Input Generation. For each dataset, we utilize the network topology and generate the input parameters as follows. The innate opinion s_i of each agent i is independently generated uniformly at random from $[0, 1]$. For each edge $\{i, j\}$ in the network, we independently pick w_{ij} uniformly at random from $[0, 1]$; otherwise, we set $w_{ij} = 0$. For $(i, j) \in V \times V$, we normalize $P_{ij} := \frac{w_{ij}}{\sum_{k \in V} w_{ik}}$. From Lemma 4.2, one can see that approximating the equilibrium opinions is more difficult when the resistance is low. However, since we still want to demonstrate that the resistance for each agent can have varied lower and upper bounds, we set the lower bound l_i of each agent i independently such that with probability 0.99, l_i equals 0.001, and with probability 0.01, it is picked uniformly at random

Algorithm 2: Optimistic Local Search

Input: Innate opinions $s \in [0, 1]^V$; interaction matrix P ; for each agent $i \in V$, upper u_i and lower l_i bounds for resistance.
Output: Optimal resistance vector $\alpha \in \times_{i \in V} \{l_i, u_i\}$.

- 1 (Technical step.) Randomly perturb each coordinate of s slightly.
- 2 Initially, for each agent i , set $\alpha_i \leftarrow u_i$ to its upper bound; denote $\epsilon_\alpha := \min_{i \in V} \alpha_i$.
- 3 Pick $z = (1, 1, \dots, 1)$, and set $t \leftarrow 0$; denote $\text{err}(t) := \frac{(1-\epsilon_\alpha)^t}{\epsilon_\alpha}$.
- 4 **while** $\exists i \in V : |s_i - z_i| \leq \text{err}(t)$ **do**
 - 5 $z \leftarrow As + (I - A)Pz$, where $A = \text{Diag}(\alpha)$
 - 6 $t \leftarrow t + 1$
 - 7 (Optimistic Candidates.) Set $L \leftarrow \{i \in V : z_i \leq s_i \wedge \alpha_i = u_i\}$.
 - 8 (Rare Mistakes.) Set $J \leftarrow \{i \in V : z_i > s_i \wedge \alpha_i = l_i\}$.
 - 9 **if** $L \cup J \neq \emptyset$ **then**
 - 10 **for each** $i \in L$ **do**
 - 11 $\alpha_i \leftarrow l_i$ to its lower bound (and update ϵ_α).
 - 12 **for each** $i \in J$ **do**
 - 13 $\alpha_i \leftarrow u_i$ to its upper bound (and update ϵ_α).
 - 14 $t \leftarrow 0$
- 15 **return** Resistance vector α .

from $[0.001, 0.1]$. Similarly, each upper bound u_i is independently selected such that with probability 0.99, it equals to 0.999, and with probability 0.01, it is chosen uniformly at random from $[0.9, 0.999]$.

5.1 Update Strategies Comparison

We compare the following three update strategies described in Section 4: conservative, opportunistic and optimistic. For the three smaller networks (com-Youtube, com-LiveJournal, LiveJournal), we apply all three update strategies. For the largest network (com-Friendster), we only report the performance of the optimistic update strategy, as the other two update strategies are not efficient enough for such a large dataset.

Experimental Setup. For fair comparison among the update strategies, we always initialize $z = (1, 1, \dots, 1)$. To compare their performances, we plot a curve for each strategy. The curves have a common x-axis, which corresponds to the number of times that the vector z has been updated so far, i.e., the number of times line 5 (in both Algorithms 1 and 2) has been executed. Since line 5 is the most time-consuming part of the algorithms, it will be a suitable common reference. We use the term *iteration* to refer to each time z is updated. For each update strategy, we explain what is plotted for the y-axis.

- **Conservative Strategy.** We run Algorithm 1 such that in line 7, L is non-empty only if for all $i \in V$, $|s_i - z_i| > \text{err}(t)$, in which case, we pick L to be the collection of all i 's such that $\alpha_i = u_i$ and $z_i \leq s_i - \text{err}(t)$. For the y-axis, we plot the ratio of agents i for which currently $\alpha_i = l_i$, or we know definitely that α_i should be switched to l_i , i.e., currently $\alpha_i = u_i$ and $z_i \leq s_i - \text{err}(t)$. In Algorithm 1, the iterations (referring to each time z is updated) are grouped into *phases*, where a non-empty L in an iteration marks the end of a phase. Observe that at the end of a phase, for all $i \in L$, α_i is set to l_i and t is reset to 0.

Table 1: Datasets Information

Name	Number n of Nodes	Number m of Edges	Description	Source
com-Youtube	1,134,890	2,987,624	Youtube online social network	[18]
com-LiveJournal	3,997,962	34,681,189	LiveJournal online social network	[18]
LiveJournal	10,690,276	112,307,385	LiveJournal user and group membership network	[17]
com-Friendster	65,608,366	1,806,067,135	Friendster online social network	[18]

Hence, in the next iteration, no coordinate α_i is certain to be switched. Hence, the curve has a step-like shape, where each plateau occurs after the end of each phase.

Observe that initially $\epsilon_\alpha = \min_i u_i \geq 0.9$. Hence, it takes very few number of iterations to satisfy $\forall i \in V : |s_i - z_i| > \text{err}(t)$; we call this the phase 0. At the end of the phase 0, we set some $\alpha_i = l_i$ and ϵ_α decreases significantly. Hence, subsequent phases have many more iterations.

Observe that we can stop the iterative process, when for all $i \in V$, $|s_i - z_i| > \text{err}(t)$, but there is no $i \in V$ such that $z_i < s_i$ and $\alpha_i = u_i$. This marks the end of the curve.

In each phase, we pick L of line 7 as the collection of all i 's such that $z_i \leq s_i - \text{err}(t)$ only when $\forall i \in V : |s_i - z_i| > \text{err}(t)$ (otherwise, we pick $L = \emptyset$). Then, we set $\alpha_i = l_i$ for each $i \in L$ and $t = 0$. We call such a phase a *conservative phase*.

- **Opportunistic Strategy.** We run Algorithm 1 similarly as before. Phase 0 is the same as the conservative strategy; we call a phase *conservative*, if it follows the conservative update strategy.

Starting from phase 1, we can perform it in an opportunistic manner as follows. Recall that at the beginning of a phase, t has just been reset to 0. At the t -th iteration of that phase, we use $L(t)$ to denote the collection of i 's such that $z_i \leq s_i - \text{err}(t)$. For every 1000 iterations, we compute an estimate slope(k) := $\frac{|L(1000k)| - |L(1000(k-1))|}{1000}$ of the slope; we keep track k_m of the maximum slope computed so far. After some iteration, if the estimated slope drops below some factor (we use 0.1 in our experiments) of k_m , then we end this phase. Intuitively, each additional iteration flips only a small number of coordinates α_i , and hence, one would like to end this phase. We call such a phase *opportunistic*.

Since typically the total number of phases is around 8, we run phase 1 to 6 opportunisticly, after which we run the remaining phases conservatively to make sure that all coordinates α_i that need to be changed will be flipped.

As in the conservative update strategy, for the y-axis, we plot the ratio of coordinates α_i that currently $\alpha_i = l_i$, or we know for sure should be switched to l_i , i.e. $\alpha_i = u_i$ and $z_i \leq s_i - \text{err}(t)$.

- **Optimistic Strategy.** We implement Algorithm 2, where in each iteration after z is updated, a coordinate α_i is (re)set to l_i if $z_i < s_i$, and (re)set to u_i if $z_i > s_i$. For the y-axis, we plot the ratio of coordinates that currently take their lower bounds. The curve ends when enough iterations are performed after some coordinate of α is last updated, in order to ensure that the estimate z is close enough to the equilibrium vector according to Lemma 4.2.




Figure 2: Update Strategies Comparison on com-Youtube.

Experiment Results. Each of Figures 2, 3 and 4 shows the plots for the three strategies in the corresponding network (com-Youtube, com-LiveJournal or LiveJournal). Figure 5 shows the plot of the optimistic strategy in the com-Friendster network, where the other two strategies are not efficient enough for such a large network. As expected, the opportunistic strategy is slightly better than the conservative strategy. From the positions of the plateaus, we can see that the initial opportunistic phases end sooner than their conservative counterparts. Hence, overall the opportunistic strategy performs slightly better than the conservative strategy; in increasing sizes of the three tested networks, the numbers of iterations taken by the opportunistic strategy are 79.2%, 77.9% and 71.5%, respectively, of those taken by the conservative strategy.

On the other hand, the optimistic strategy can achieve the optimal resistance vector with much fewer number of iterations than the other two strategies. In increasing sizes of the three smaller networks, the numbers of iterations taken by the optimistic strategy are only 12.8%, 13.4% and 12.4%, respectively, of those taken by the conservative strategy. Moreover, the optimistic strategy makes very few mistakes; in increasing sizes of the four networks, the numbers of times coordinates are flipped from lower bounds back to upper bounds are 1, 0, 13 and 168, which are negligible for networks with millions of nodes.

5.2 Running Time with Multiple Threads

We compare the actual running time using different number of threads for the optimistic strategy on only the three smaller networks, since the largest network takes too long using only one thread. Using all 24 available threads, running the optimistic strategy on the com-Friendster network already takes around 50 hours.

The three bar graphs in Figure 6 show the running time (measured in minutes) for running the optimistic strategy with different




Figure 3: Update Strategies Comparison on com-LiveJournal.




Figure 4: Update Strategies Comparison on LiveJournal.




Figure 5: Optimistic Strategy on com-Friendster.

number of threads on the com-YouTube, com-LiveJournal and LiveJournal networks. Since updating z (line 5 of Algorithm 2) is the most time-consuming part of the algorithm, the fact that it is readily parallelizable supports the empirical results that using multiple threads can greatly reduce the running time, where the effect is more prominent for larger networks.

6 TECHNICAL PROOFS

LEMMA 6.1. (Lemma 3.1 restated) Given $K \subsetneq V$ and $\alpha \in (0, 1)^V$, let $A := \text{Diag}(\alpha)$ and recall that P is the irreducible interaction matrix. Then, the inverse $M = [I - (I - A_{-K})P]^{-1}$ exists, and every entry




Figure 6: Running Time with Different Number of Threads on com-YouTube, com-LiveJournal and LiveJournal

of M is positive. Moreover, for each $k \in V$, define $a_k = 0$ if $k \in K$, otherwise $a_k = \alpha_k$. Then, we have:

- (1) $(PM)_{kk} = \frac{M_{kk}-1}{1-a_k} > 0$;
- (2) $(PM)_{kj} = \frac{M_{kj}}{1-a_k} > 0$, for each $j \neq k$.

PROOF. Observe that P corresponds to an irreducible random walk. Hence, $(I - A_{-K})P$ represents a diluted random walk, where at the beginning of each step, the measure at nodes $i \notin K$ will suffer a factor of $1 - \alpha_i \in (0, 1)$. The irreducibility of the random walk P means that every state is reachable from any state. Hence, starting from any measure vector, eventually the measure at every node will tend to 0. This means that $(I - A_{-K})P$ has eigenvalues with magnitude strictly less than 1. Therefore, we can consider the following Neumann series of a matrix:

$$M = [I - (I - A_{-K})P]^{-1} = I + \sum_{k=1}^{\infty} [(I - A_{-K})P]^k,$$

which implies that the inverse M exists, and every entry of M is positive; in particular, for every $k \in V$, $M_{kk} > 1$.

By the definition of M , we have $[I - (I - A_{-K})P]M = I$. We fix some $k \in V$. By considering the (k, k) -the entry, i.e., the dot product between the k -th row of $[I - (I - A_{-K})P]$ and the k -th column of M , we have

$$M_{kk} - \sum_{i \in V} (1 - a_k)P_{ki}M_{ik} = 1$$

where $a_k = 0$ if $k \in K$, otherwise $a_k = \alpha_k < 1$. Hence, we have $(PM)_{kk} = \sum_{i \in V} P_{ki}M_{ik} = \frac{M_{kk}-1}{1-a_k}$.

Similarly, for $j \neq k$, by considering the dot product between the k -th row of $[I - (I - A_{-K})P]$ and the j -th column of M , we have

$$M_{kj} - \sum_{i \in V} (1 - a_k)P_{ki}M_{ij} = 0.$$

Hence, we have for $j \neq k$,

$$(PM)_{kj} = \sum_{i \in V} P_{ki}M_{ij} = \frac{M_{kj}}{1-a_k},$$

as required. \square

LEMMA 6.2. (Lemma 3.3 restated) Referring to Lemma 3.2. For any $\alpha \in (0, 1)^V$ and any $i \in V$, denote $M = [I - (I - A_{-\{i\}})P]^{-1}$. Then, $\frac{\partial f(\alpha)}{\partial \alpha_i}$ has the same sign in $\{-, 0, +\}$ as $s_i - \sum_{j \neq i} M_{ij}\alpha_j s_j$, which is independent of α_i .

PROOF. Using the Sherman-Morrison formula, we consider

$$\begin{aligned} X &:= [I - (I - A)P]^{-1} = [I - (I - A_{-\{i\}} - \alpha_i \mathbf{e}_i \mathbf{e}_i^T)P]^{-1} \\ &= [I - (I - A_{-\{i\}})P + \alpha_i \mathbf{e}_i \mathbf{e}_i^T P]^{-1} \\ &= M - \frac{\alpha_i}{1 + \alpha_i \mathbf{e}_i^T P M \mathbf{e}_i} M \mathbf{e}_i \mathbf{e}_i^T P M. \end{aligned}$$

Observe that $\mathbf{e}_i^T P M \mathbf{e}_i = (PM)_{ii}$ and $(M \mathbf{e}_i \mathbf{e}_i^T P M)_{ij} = M_{ii}(PM)_{ij}$ for each $j \in V$. Then, by Lemma 3.1 with $K = \{i\}$, we have

$$\begin{aligned} X_{ii} &= M_{ii} - \frac{\alpha_i M_{ii}(PM)_{ii}}{1 + \alpha_i (PM)_{ii}} = M_{ii} - \frac{\alpha_i M_{ii}(M_{ii} - 1)}{1 + \alpha_i (M_{ii} - 1)} \\ &= \frac{(1 - \alpha_i + \alpha_i M_{ii} - \alpha_i M_{ii} + \alpha_i) M_{ii}}{1 - \alpha_i + \alpha_i M_{ii}} = \frac{M_{ii}}{1 - \alpha_i + \alpha_i M_{ii}}; \end{aligned}$$

and for $j \neq i$,

$$\begin{aligned} X_{ij} &= M_{ij} - \frac{\alpha_i M_{ii}(PM)_{ij}}{1 + \alpha_i (PM)_{ii}} = M_{ij} - \frac{\alpha_i M_{ii} M_{ij}}{1 + \alpha_i (M_{ii} - 1)} \\ &= \frac{(1 - \alpha_i + \alpha_i M_{ii} - \alpha_i M_{ii}) M_{ij}}{1 - \alpha_i + \alpha_i M_{ii}} = \frac{(1 - \alpha_i) M_{ij}}{1 - \alpha_i + \alpha_i M_{ii}}. \end{aligned}$$

By Lemma 3.2, we know $\frac{\partial f(\alpha)}{\partial \alpha_i}$ and $s_i - z_i(\alpha)$ have the same sign in $\{-, 0, +\}$. Recall that $z(\alpha) = [I - (I - A)P]^{-1} A s = X A s$. Applying the above results, we have

$$\begin{aligned} s_i - z_i(\alpha) &= s_i - \sum_{j \in V} X_{ij} \alpha_j s_j \\ &= s_i - \frac{M_{ii} \alpha_i s_i}{1 - \alpha_i + \alpha_i M_{ii}} - \sum_{j \neq i} \frac{(1 - \alpha_i) M_{ij} \alpha_j s_j}{1 - \alpha_i + \alpha_i M_{ii}} \\ &= \frac{s_i (1 - \alpha_i + \alpha_i M_{ii}) - M_{ii} \alpha_i s_i - (1 - \alpha_i) \sum_{j \neq i} M_{ij} \alpha_j s_j}{1 - \alpha_i + \alpha_i M_{ii}} \\ &= \frac{(1 - \alpha_i) (s_i - \sum_{j \neq i} M_{ij} \alpha_j s_j)}{1 - \alpha_i + \alpha_i M_{ii}}. \end{aligned}$$

Since $\alpha_i \in (0, 1)$, we conclude that $\frac{1 - \alpha_i}{1 - \alpha_i + \alpha_i M_{ii}} > 0$. Thus $\frac{\partial f(\alpha)}{\partial \alpha_i}$, $s_i - z_i(\alpha)$ and $s_i - \sum_{j \neq i} M_{ij} \alpha_j s_j$ have the same sign. \square

LEMMA 6.3. (Lemma 3.6 restated) For any $i, k \in V$ such that $i \neq k$, let $M = [I - (I - A_{-\{i\}})P]^{-1}$ and $R = [I - (I - A_{-\{i, k\}})P]^{-1}$. Then for any $j \in V$, we have

- (1) $M_{jk} = \frac{R_{jk}}{1 + \alpha_k R_{kk} - \alpha_k}$,
- (2) $M_{jh} = R_{jh} - \frac{\alpha_k R_{jk} R_{kh}}{1 + \alpha_k R_{kk} - \alpha_k}$, for each $h \neq k$.

In particular, the quantity in Lemma 3.3 can be rewritten as follows:

$$s_i - \sum_{j \neq i} M_{ij} \alpha_j s_j = s_i - \sum_{j \neq i, k} R_{ij} \alpha_j s_j - \frac{\alpha_k R_{ik}}{1 + \alpha_k R_{kk} - \alpha_k} (s_k - \sum_{j \neq i, k} R_{kj} \alpha_j s_j).$$

PROOF. Using the Sherman-Morrison formula, we have

$$\begin{aligned} M &= [I - (I - A_{-\{i, k\}})P + \alpha_k \mathbf{e}_k \mathbf{e}_k^T P]^{-1} \\ &= R - \frac{\alpha_k}{1 + \alpha_k \mathbf{e}_k^T P R \mathbf{e}_k} R \mathbf{e}_k \mathbf{e}_k^T P R \end{aligned}$$

We can compute that $\mathbf{e}_k^T P R \mathbf{e}_k = (PR)_{kk}$ and $(R \mathbf{e}_k \mathbf{e}_k^T P R)_{jh} = R_{jk}(PR)_{kh}$ for $j, h \in V$. Then we have

$$M_{jh} = R_{jh} - \frac{\alpha_k R_{jk}(PR)_{kh}}{1 + \alpha_k (PR)_{kk}}.$$

By Lemma 3.1, we obtain

$$M_{jh} = R_{jh} - \frac{\alpha_k R_{jk} R_{kh}}{1 + \alpha_k R_{kk} - \alpha_k} \text{ for } j, h \in V \text{ and } h \neq k,$$

and

$$M_{jk} = R_{jk} - \frac{\alpha_k R_{jk}(R_{kk} - 1)}{1 + \alpha_k R_{kk} - \alpha_k} = \frac{R_{jk}}{1 + \alpha_k R_{kk} - \alpha_k} \text{ for } j \in V.$$

as required. \square

LEMMA 6.4. (Lemma 3.7 restated) Suppose $\alpha \in (0, 1)^V$, recall that $A_{-\{i, k\}} := \text{Diag}(\alpha_{-\{i, k\}})$, and P corresponds to an irreducible interaction matrix. For any $i, k \in V$ such that $i \neq k$, let $R = [I - (I - A_{-\{i, k\}})P]^{-1}$, then $R_{ii} = \max_{j \in V} R_{ji}$. Moreover, $R_{ii} = R_{ki}$ if and only if $P_{kk} + P_{ki} = 1$.

PROOF. We have $[I - (I - A_{-\{i, k\}})P]R = I$. By considering the dot product between each row of $[I - (I - A_{-\{i, k\}})P]$ and column i of R , we have

$$R_{ii} - \sum_{h \in V} P_{ih} R_{hi} = 1, R_{ki} - \sum_{h \in V} P_{kh} R_{hi} = 0,$$

$$\text{and } R_{ji} - \sum_{h \in V} (1 - \alpha_j) P_{jh} R_{hi} = 0, \text{ for } j \neq i, k.$$

After rearranging, we have

$$R_{ii} = 1 + \sum_{h \in V} P_{ih} R_{hi}, R_{ki} = \sum_{h \in V} P_{kh} R_{hi},$$

$$\text{and } R_{ji} = (1 - \alpha_j) \sum_{h \in V} P_{jh} R_{hi}, \text{ for } j \neq i, k.$$

Now it suffices to show that for $j \neq i, k$, the above R_{ji} cannot be the maximum among them, and $R_{ki} \leq R_{ii}$.

First, we show that R_{ji} cannot be the maximum. Since $\sum_{h \in V} P_{jh} = 1$ and $\alpha_j \in (0, 1)$, we have

$$R_{ji} = (1 - \alpha_j) \sum_{h \in V} P_{jh} R_{hi} \leq (1 - \alpha_j) \max_{h \in V} R_{hi} < \max_{h \in V} R_{hi}.$$

Thus, R_{ji} cannot be the maximum.

Next, we show that $R_{ki} \leq R_{ii}$ by contradiction. Suppose $R_{ki} > R_{ii}$, then R_{ki} is the unique maximum in the i -th column of R . Since $\sum_{h \in V} P_{kh} = 1$ and $R_{ki} = \sum_{h \in V} P_{kh} R_{hi}$, it must be the case that $P_{kk} = 1$. This means P corresponds to a random walk with absorbing state k , which contradicts that P is irreducible. Therefore, we have $R_{ki} \leq R_{ii}$, and hence $R_{ii} = \max_{h \in V} R_{hi}$.

Observe that we already know $R_{ji} < R_{ii}$ for $j \neq i, k$, and $R_{ki} = \sum_{h \in V} P_{kh} R_{hi}$. Hence, $R_{ki} = R_{ii}$ implies that $P_{kk} + P_{ki} = 1$.

Conversely, $P_{kk} + P_{ki} = 1$ implies that $R_{ki} = P_{kk} R_{ki} + P_{ki} R_{ii}$. As argued above, we must have $P_{kk} \neq 1$, which implies $R_{ki} = R_{ii}$. \square

7 CONCLUSION AND FUTURE WORK

We have given a solid theoretical analysis of the unbudgeted variant of the opinion susceptibility problem, and designed scalable local search algorithms that can solve the problem optimally on graphs with millions of nodes.

We believe that our techniques will lead to insights for the analysis of the budgeted variant of the problem. We leave the task of providing theoretical guarantees for greedy algorithms on the budgeted variant as future work.

REFERENCES

- [1] Rediet Abebe, Jon M. Kleinberg, David C. Parkes, and Charalampos E. Tsourakakis. 2018. Opinion Dynamics with Varying Susceptibility to Persuasion. In *KDD*. ACM, 1089–1098.
- [2] Joseph P. Allen, Maryfrances R. Porter, and F. Christy McFarland. 2006. Leaders and followers in adolescent close friendships: Susceptibility to peer influence as a predictor of risky behavior, friendship instability, and depression. *Development and psychopathology* 18, 1 (2006), 155–172.
- [3] Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. 2016. Global Optimality of Local Search for Low Rank Matrix Recovery. In *NIPS*. 3873–3881.
- [4] Morris H. DeGroot. 1974. Reaching a Consensus. *J. Amer. Statist. Assoc.* 69, 345 (1974), 118–121. <http://www.jstor.org/stable/2285509>
- [5] T. E. Dielman, Deborah D. Kloska, Sharon L. Leech, John E. Schulenberg, and Jean T. Shope. 1992. Susceptibility to peer pressure as an explanatory variable for the differential effectiveness of an alcohol misuse prevention program in elementary schools. *Journal of School Health* 62, 6 (1992), 233–237.
- [6] Pedro Domingos and Matt Richardson. 2001. Mining the Network Value of Customers. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '01)*. ACM, New York, NY, USA, 57–66. <https://doi.org/10.1145/502512.502525>
- [7] William N. Evans, Wallace E. Oates, and Robert M. Schwab. 1992. Measuring peer group effects: A study of teenage behavior. *Journal of Political Economy* 100, 5 (1992), 966–991.
- [8] Noah E. Friedkin and Eugene C. Johnsen. 1999. Social Influence Networks and Opinion Change. 16 (01 1999), 1–19.
- [9] Rong Ge, Furong Huang, Chi Jin, and Yang Yuan. 2015. Escaping From Saddle Points - Online Stochastic Gradient for Tensor Decomposition. In *COLT (JMLR Workshop and Conference Proceedings)*, Vol. 40. JMLR.org, 797–842.
- [10] Rong Ge, Jason D. Lee, and Tengyu Ma. 2016. Matrix Completion has No Spurious Local Minimum. In *NIPS*. 2973–2981.
- [11] Aristides Gionis, Evimaria Terzi, and Panayiotis Tsaparas. 2013. Opinion Maximization in Social Networks. In *SDM*. SIAM, 387–395.
- [12] Philip Hartman. 1959. On functions representable as a difference of convex functions. *Pacific J. Math.* 9, 3 (1959), 707–713. <https://projecteuclid.org:443/euclid.pjm/1103039111>
- [13] Prateek Jain and Purushottam Kar. 2017. Non-convex Optimization for Machine Learning. *Foundations and Trends in Machine Learning* 10, 3-4 (2017), 142–336.
- [14] Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M. Kakade, and Michael I. Jordan. 2017. How to Escape Saddle Points Efficiently. In *ICML (Proceedings of Machine Learning Research)*, Vol. 70. PMLR, 1724–1732.
- [15] Maurits Kaptein, Panos Markopoulos, Boris Ruyter, and Emile Aarts. 2009. Can You Be Persuaded? Individual Differences in Susceptibility to Persuasion. In *Proceedings of the 12th IFIP TC 13 International Conference on Human-Computer Interaction: Part I (INTERACT '09)*. Springer-Verlag, Berlin, Heidelberg, 115–118. https://doi.org/10.1007/978-3-642-03655-2_13
- [16] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the Spread of Influence Through a Social Network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '03)*. ACM, New York, NY, USA, 137–146. <https://doi.org/10.1145/956750.956769>
- [17] Jérôme Kunegis. 2017. KONECT: Konect Network Dataset. <http://konect.uni-koblenz.de>.
- [18] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>.
- [19] Kevin M. Murphy and Andrei Shleifer. 2004. Persuasion in politics. *American Economic Review* 94, 2 (2004), 435–439.
- [20] Dohyung Park, Anastasios Kyrillidis, Constantine Caramanis, and Sujay Sanghavi. 2017. Non-square matrix sensing without spurious local minima via the Burer-Monteiro approach. In *AISTATS (Proceedings of Machine Learning Research)*, Vol. 54. PMLR, 65–74.
- [21] Ju Sun, Qing Qu, and John Wright. 2016. A geometric analysis of phase retrieval. In *ISIT*. IEEE, 2379–2383.
- [22] Ju Sun, Qing Qu, and John Wright. 2017. Complete Dictionary Recovery Over the Sphere I: Overview and the Geometric Picture. *IEEE Trans. Information Theory* 63, 2 (2017), 853–884.