Normalized Friedkin–Johnsen Opinion Dynamics

Runze Zhang, Gengyu Wang, Zhongzhi Zhang*
College of Computer Science and Artificial Intelligence
Fudan University
{rzzhang24, gywang24}@m.fudan.edu.cn, zhangzz@fudan.edu.cn

Abstract

As online social platforms increasingly serve as central arenas for opinion aggregation, numerous opinion dynamics models have been extensively developed to characterize how opinions propagate and to predict emerging social phenomena. Among various models, the Friedkin–Johnsen (FJ) model stands out as one of the most extensively studied frameworks, with broad applications across diverse domains and recent inspiration on the design of graph neural network architectures. However, the conventional FJ row-stochastic pairwise interaction has been shown to induce imbalanced influence patterns, often weakening the impact of high-degree nodes. To address this issue, we propose a normalized FJ model based on the normalized Laplacian, which symmetrically incorporates degree information from both endpoints. We establish its convergence and analyze its structural properties and intrinsic balance, along with a game-theoretic interpretation. We further develop efficient algorithms for computing equilibrium opinions and related metrics. Experiments on large-scale networks demonstrate the effectiveness and scalability of our approach.

1 Introduction

With the rapid proliferation of online platforms, social interactions have become increasingly interconnected and instantaneous, fundamentally reshaping how information and opinions spread [31]. In such environments, individual viewpoints are no longer formed in isolation but are continuously influenced and adjusted through repeated interactions within large-scale online social networks [37]. The intricate collective patterns arising from these interactions, commonly referred to as *opinion dynamics*, provide valuable insights into social behavior, yet their underlying mechanisms remain not fully understood [4, 51]. From a system-theoretic perspective, this motivates the development of mathematical models for opinion aggregation that strike a balance between expressive power and analytical tractability [15, 2, 17, 23]. Among the existing models, the Friedkin–Johnsen (FJ) framework [17] stands out for its simplicity and flexibility in capturing realistic opinion evolution and its proven alignment with real-world phenomena [55]. Extending the classical DeGroot model [15], FJ incorporates persistent intrinsic opinions for each individual, leading to a linear discrete-time dynamical system in which expressed opinions evolve under social influence while remaining anchored to internal beliefs. This mechanism allows the model to capture a rich set of phenomena beyond consensus, including disagreement [35], conflict [12], and polarization [35, 34], thereby providing a more faithful representation of real-world social dynamics.

Research on the FJ model spans several interconnected directions. A central line of work focuses on efficiently computing the equilibrium opinions induced by the dynamics [53, 36]. Another important direction studies opinion optimization, aiming to steer the equilibrium outcome under budget constraints by intervening on selected individuals or the network structure. Typical approaches

*Corresponding author.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

include modifying innate opinions [57], expressed opinions [20], or susceptibility to persuasion [2, 10], as well as altering the network topology through operations such as edge addition [66, 68]. In parallel, substantial efforts have been devoted to extending the original model to enhance its expressive power, including incorporating external influences such as biased media via fixed-opinion agents [38], modeling antagonistic relationships through negatively weighted edges [65], and generalizing to multidimensional settings that capture interdependencies across multiple topics [39].

It is worth mentioning that, beyond its role in modeling opinion dynamics, the FJ model has recently emerged as a unifying perspective bridging computational social science and graph-based machine learning. In particular, the message passing paradigm in graph neural networks can be interpreted through the lens of FJ dynamics, where residual connections naturally encode the trade-off between social influence and individual stubbornness [47]. This connection further highlights the broad applicability of the FJ model and underscores the importance of studying it in greater depth.

So far, most studies on FJ opinion dynamics have focused on the standard model [50, 58], which is most thoroughly described and analyzed in [9]. However, this model inherently implies that individuals with more neighbors place less weight on their innate opinions, a property that has been pointed out to be unrealistic under certain conditions [11]. To address this limitation, a normalized extension of the original FJ model is preferred. Notably, Friedkin and Johnsen themselves suggested that interpersonal influence should be normalized to better capture realistic social interactions [18]. Despite some previous efforts [11, 2], existing normalization strategies remain insufficient to fully capture the bilateral nature of social interactions, as they typically depend on the local structure of a single endpoint. This observation motivates our study: we seek to develop a principled normalization framework that simultaneously accounts for the structural characteristics of both individuals involved in an interaction, thereby providing a more balanced and realistic modeling of influence.

In this work, we propose a normalized variant of the classical FJ model based on a principled treatment of interpersonal influence via the normalized Laplacian, where pairwise interactions incorporate the degree information of both endpoints. This formulation yields a more symmetric and structurally consistent characterization of social influence, while preserving the key features of FJ-type dynamics and inducing a new equilibrium behavior. We establish the convergence of the proposed model and analyze its structural properties and intrinsic balance, complemented by a game-theoretic interpretation. Furthermore, the proposed formulation reveals connections to feature propagation in graph convolutional networks [56] and label spreading in semi-supervised learning [63]. To enable practical deployment, we develop efficient algorithms for computing equilibrium opinions and related descriptive metrics. Extensive experiments on large-scale networks demonstrate the effectiveness and scalability of our approach.

2 Related Works

To understand and model the spread of public opinion in social networks, researchers have proposed a variety of mathematical models, including the DeGroot model [15], the Hegselmann–Krause (HK) model [23], the voter model [25], and many others [27, 40, 16, 5]. Among these, the Friedkin–Johnsen (FJ) model [17] has attracted significant attention due to its simplicity, extensibility, and strong empirical relevance [55, 8] and has been widely adopted in tasks such as opinion computation [8, 53, 36, 58], public-opinion optimization [20, 50, 65, 54, 68], and adversarial defense for network security [67, 69]. Beyond simple graphs, studies of the FJ model have also been extended to directed graphs [50], signed graphs [65], time-varying graphs [41], higher-order relations [60], long-range neighbors [61], noisy networks [59], and multidimensional networks [39].

Mathematically, the FJ opinion evolution formula can be interpreted as an absorbing random walk process [20] and is closely related to PageRank [36]. The coefficient matrix of the equilibrium opinion corresponds to a diagonally shifted Laplacian matrix, which is commonly studied in computer graphics [29], and its entries admit meaningful interpretations via the matrix forest theorem [6, 50]. Besides traditional mathematical analysis, recent studies also leverage graph neural networks (GNNs) to learn and approximate the evolution process of FJ opinion dynamics [32]. Conversely, a recent line of work in GNNs further draws inspiration from FJ opinion dynamics by incorporating such diagonal heterogeneous resistance parameters [2] into the control of information propagation. In particular, [47] introduces an adaptive initial residual connection mechanism, where node-specific residual strengths

balance initial features and neighborhood aggregation, thereby effectively mitigating oversmoothing and achieving strong empirical performance.

A complementary line of research examines uncertainty in opinion dynamics. Stochastic gossip models with stubborn agents explain persistent opinion fluctuations and disagreement [3], while noisy extensions of DeGroot and FJ dynamics characterize how network topology affects opinion diversity [49]. Within the FJ framework, probabilistic innate opinions have been used to select informative observations for estimating collective opinions [43], and limited access to innate opinions motivates intervention methods that account for reconstruction errors [13]. These studies highlight the roles of stochastic interactions, uncertain inputs, and partial observation in shaping or estimating collective outcomes.

Despite the wide range of applications of the FJ model across various domains, the issue of its proper normalization remains an important topic for further investigation. The works most closely related to this topic are [11, 2], which introduce homogeneous or heterogeneous stubbornness parameters to rebalance the trade-off between individuals' intrinsic conflicts and external social influence. However, despite incorporating heterogeneity [2], this formulation still relies on row-normalized weights and therefore fails to properly capture the strength of pairwise interactions, as it depends solely on the degree of one endpoint rather than the joint structural properties of node pairs.

3 Preliminaries

Notations. We use lowercase letters such as a, b, and c to represent scalars in $\mathbb{R}$, uppercase letters such as A, B, and C to denote sets, bold lowercase letters such as $\mathbf{a}$, $\mathbf{b}$, and $\mathbf{c}$ to represent vectors, and bold uppercase letters such as $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ to represent matrices. The transpose of a vector $\mathbf{a}$ and a matrix $\mathbf{A}$ are denoted by $\mathbf{a}^\top$ and $\mathbf{A}^\top$, respectively. The (i,j)-th entry of a matrix $\mathbf{A}$ is written as $a_{i,j}$ or $\mathbf{A}[i,j]$, while the i-th entry of a vector $\mathbf{a}$ is denoted as a_i or $\mathbf{a}[i]$. The spectral radius $\rho(\mathbf{A})$ of matrix $\mathbf{A}$ is the maximum modulus of its eigenvalues, i.e., $\rho(\mathbf{A}) = \max_{i\in\{1,2,\cdots,n\}} |\lambda_i|$, where n is the dimension of $\mathbf{A}$. All proofs of the theorems stated in this paper are provided in the appendix.

Graph. Following many related studies[2, 36, 65], we denote a social network by a connected, undirected and weighted graph $\mathcal{G} = (V, E)$ with $n = |V|$ nodes and $m = |E|$ edges, where $V = \{v_1, v_2, \ldots, v_n\}$ is the set of nodes, and $E \subseteq \{\{v_i, v_j\} \in V \times V\}$ is the set of edges. Let $\mathbf{W} = [w_{i,j}] \in \mathbb{R}^{n\times n}$ denote the weighted adjacency matrix of $\mathcal{G}$, where $w_{i,j} > 0$ if $\{v_i, v_j\} \in E$, and $w_{i,j} = 0$ otherwise. Since $\mathcal{G}$ is undirected, $\mathbf{W}$ is symmetric, i.e., $\mathbf{W} = \mathbf{W}^\top$. The degree of node v_i is defined as $d_i = \sum_{j=1}^{n} w_{i,j}$, and $\mathbf{D} = \mathrm{diag}(d_1, d_2, \ldots, d_n)$ denotes the diagonal degree matrix. The Laplacian matrix of $\mathcal{G}$ is defined to be $\mathbf{L} = \mathbf{D} - \mathbf{W}$. The symmetric normalized Laplacian matrix $\mathbf{L}_{\text{sym}}$ is defined as $\mathbf{L}_{\text{sym}} = \mathbf{D}^{-\frac{1}{2}}\mathbf{L}\mathbf{D}^{-\frac{1}{2}} = \mathbf{I} - \mathbf{D}^{-\frac{1}{2}}\mathbf{W}\mathbf{D}^{-\frac{1}{2}}$.

Vanilla FJ. Under the FJ framework [17], each individual in a social network, represented by a node in graph $\mathcal{G} = (V, E)$, is assigned two opinion values. One is the *innate opinion*, denoted as s_i, which reflects the unchangeable belief of node i on a particular topic and remains constant throughout the entire dynamic process. The other is the *expressed opinion*, denoted as $z_i(t)$, which is shared with its neighbors and evolves over time, being in the same domain as s_i. The range of the opinion is typically chosen as the interval $[0, 1]$, where 0 and 1 represent the two extreme views on a given topic.

Each agent is supposed to update its expressed opinion through weighted averaging its innate opinion and the expressed opinions of all its out neighbors at the previous time step, following the equation below:

$$z_i(t+1) = \frac{s_i + \sum_{j\in N_i} w_{i,j} z_j(t)}{1 + d_i}. \tag{1}$$

It is proven that this opinion evolution eventually admits a unique Nash equilibrium [9] . Further, let us define the initial opinion vector as $\mathbf{s} = (s_1, s_2, \ldots, s_n)^\top$, and define the expressed opinion vector at time t as $\mathbf{z}(t) = (z_1(t), z_2(t), \ldots, z_n(t))^\top$. As t approaches infinity, the expressed opinion vector converges to an equilibrium opinion vector $\mathbf{z}^* = (z_1, z_2, \ldots, z_n)^\top$, satisfying:

$$\mathbf{z}^* = \lim_{t\to\infty} \mathbf{z}(t) = (\mathbf{I} + \mathbf{L})^{-1}\mathbf{s}. \tag{2}$$

From a game-theoretic perspective, the FJ opinion dynamics can be derived from a social opinion game [9]. It is assumed that each individual updates their expressed opinions in order to minimize

his cognitive dissonance. More precisely, at each time step, node i updates its expressed opinion by attempting to minimize its psycho-social cost [21] defined below:

$$c(z_i) \triangleq (z_i - s_i)^2 + \sum_{j \in N_i} w_{i,j}(z_i - z_j)^2, \tag{3}$$

which is a function that measures both the difference between an individual's expressed opinion and innate opinion, as well as the total difference between the expressed opinions of its outgoing neighbors and its own, weighted by social influence [9, 42, 22]. These two components of the psycho-social cost function are commonly referred to as *internal conflict* and *external conflict* in existing research [65, 12]. By taking the derivative of (3) with respect to the decision variable and setting it to zero, we obtain (1).

A pioneering attempt to normalize the FJ model was proposed in [11], where an additional parameter is introduced to balance the relative importance of internal and external conflicts in the cost function. Specifically, [11] considers the following opinion update rule:

$$z_i(t+1) = \frac{a}{a+1} s_i + \frac{1}{a+1} \sum_{j \in N_i} \frac{w_{i,j}}{d_i} z_j(t),$$

which controls the extent to which individuals rely on their intrinsic opinions versus social influence through a single global parameter. This formulation can be viewed as a homogeneous instance of the more general FJ model with heterogeneous resistance parameters [2], given by

$$z_i(t+1) = \alpha_i s_i + (1 - \alpha_i) \sum_{j \in N_i} \frac{w_{i,j}}{d_i} z_j(t), \tag{4}$$

where each individual is associated with a distinct resistance parameter α_i.

4 Normalized Friedkin-Johnsen Model

4.1 Model Definition

A central challenge in modeling opinion dynamics lies in how to properly normalize pairwise influence between individuals. A conventional approach is to employ a row-stochastic matrix to model normalized pairwise interactions. As shown in (1) and (4), the edge weight $w_{i,j}$ is normalized by $d_i + 1$ and d_i, respectively. However, such normalization relies solely on the degree of a single endpoint, thereby overlooking the inherently mutual nature of pairwise interactions and potentially leading to a biased characterization of influence strength. Empirical evidence from real-world networks further underscores the limitations of this one-sided treatment. In particular, [26] shows that, compared to single-degree or additive formulations, the degree product $d_i d_j$ exhibits a stronger correlation with edge betweenness centrality. Moreover, [33] demonstrates that the average interaction strength between nodes scales proportionally to $(d_i d_j)^\theta$, where a commonly observed value is $\theta = 0.5$. This pattern is further supported by [7], suggesting that interaction strength is fundamentally a joint property of both endpoints rather than being determined in isolation.

Motivated by these findings, we adopt $\frac{w_{i,j}}{\sqrt{d_i d_j}}$ as a principled normalization for pairwise interaction weights. Meanwhile, following [26, 2], we introduce an individual-specific parameter α_i to balance the influence between inner conflict and social pressure. The resulting opinion dynamics can be expressed in both node-wise and matrix forms as

$$z_i(t+1) = \alpha_i s_i + (1 - \alpha_i) \sum_{j \in N_i} w_{i,j} \frac{z_j(t)}{\sqrt{d_i d_j}} \quad \text{and} \quad \mathbf{z}(t+1) = \mathbf{D}_\alpha \mathbf{s} + (\mathbf{I} - \mathbf{D}_\alpha)\mathbf{S}\mathbf{z}(t), \tag{5}$$

where $\mathbf{z}(t) = [z_1(t), z_2(t) \ldots, z_n(t)]^T$, $\mathbf{s} = [s_1, s_2, \ldots, s_n]^T$, $\mathbf{D}_\alpha = \mathrm{diag}(\alpha_1, \alpha_2, \ldots, \alpha_n)$, and $\mathbf{S} = \mathbf{D}^{-1/2}\mathbf{W}\mathbf{D}^{-1/2}$. We hereafter refer to α_i as the resistance parameter. Following the same setting in [2], we assume that $\alpha_i \in [\alpha_{\min}, \alpha_{\max}] \subset (0, 1)$ for all $i \in V$.

Symmetry of the normalized pairwise weights is compatible with heterogeneous individual responses. Effective influence also depends on the recipient's resistance, so two nodes need not exert equal influence on each other.

The convergence of the proposed model is characterized by the following theorem.

Theorem 1. *As $t \to \infty$, the dynamics in* (5) *converges to a unique equilibrium, given by*

$$\mathbf{z}^* \triangleq \lim_{t\to\infty} \mathbf{z}(t) = \left(\mathbf{I} - (\mathbf{I} - \mathbf{D}_\alpha)\mathbf{S}\right)^{-1}\mathbf{D}_\alpha \mathbf{s}. \tag{6}$$

Compared with the most closely related formulations [9, 11, 2], the proposed model exhibits several notable advantages. From a practical perspective, it provides a more realistic and meaningful characterization of pairwise interactions between individuals. From a mathematical standpoint, the resulting interaction matrix $\mathbf{S}$, in contrast to the commonly used row-stochastic matrix $\mathbf{P} = \mathbf{D}^{-1}\mathbf{W}$, is symmetric and thus admits an orthogonal eigendecomposition. In addition, an important distinction between the proposed model and the classical model [9] is that the sum of opinions is no longer conserved. In (6), the coefficient matrix associated with the initial opinion vector $\mathbf{s}$ is no longer row-stochastic or strictly diagonally dominant. As a result, it cannot be directly interpreted as the transition matrix of an absorbing random walk, and several commonly used computational tools [6, 48] are therefore not directly applicable.

A simplified version of our dynamic can be derived by adopting the same assumption as in [9], namely that individuals place equal emphasis on internal and external conflicts. Specifically, we set $\forall i \in V, \alpha_i = \frac{1}{2}$. Following the same derivation as above, it is straightforward to obtain

$$z_i(t+1) = \frac{1}{2}s_i + \frac{1}{2}\sum_{j\in N_i} \frac{w_{i,j}}{\sqrt{d_i d_j}} z_j(t), \quad \text{and} \quad \mathbf{z}^* = (\mathbf{I} + \mathbf{L}_{\text{sym}})^{-1}\mathbf{s}, \tag{7}$$

which is consistent with formulation (2) commonly used in prior studies. We do not treat (7) as a standalone model in this paper. Nevertheless, it is worth noting that it naturally arises as a special case and is compatible with all subsequent properties and algorithms developed in this work.

4.2 Game-Theoretic Interpretation and Cross-Domain Insights

The derived opinion evolution process in (5) is not merely a mathematical reformulation, but also admits a principled game-theoretic interpretation. Specifically, we define the individual cost function as

$$c_i(z_i) \triangleq \frac{\alpha_i}{1-\alpha_i}\|z_i - s_i\|_2^2 + \sum_{j\in N_i} w_{i,j} \left\| \frac{z_i}{\sqrt{d_i}} - \frac{z_j}{\sqrt{d_j}} \right\|_2^2, \tag{8}$$

which extends the vanilla formulation in (3) in two key aspects. First, the parameter α_i captures heterogeneous resistance levels across individuals, and its equivalent reweighting $\frac{\alpha_i}{1-\alpha_i}$ aligns our formulation with the general opinion game with arbitrary node weights studied in [9], which ensures that our model is grounded in a well-established framework. Second, the pairwise disagreement term is symmetrically normalized, leading to a more principled weighting of social interactions, which corresponds to the Dirichlet energy [64] of the node representations over the graph and is widely adopted in graph learning and message passing neural networks to enforce smoothness across adjacent nodes [62]. Taking the derivative of c_i with respect to z_i and applying a discrete-time update scheme, directly leads to the update rule in (5).

In fact, this form of modeling pairwise interaction has also been extensively justified and applied in the fields of semi-supervised learning and graph neural networks. Specifically, in [63], the second term in (8) is referred to as a *smoothness constraint*, which enforces that a desirable classification function should vary smoothly over neighboring nodes. Similarly, in [56], a related formulation is employed at each layer of a graph convolutional network to propagate and average feature representations over local neighborhoods. To the best of our knowledge, our work is the first to incorporate such principles and mechanisms into opinion dynamics. Innate opinions act as persistent node inputs, while individual resistance controls their retention during neighborhood propagation. This correspondence connects the analyzed dynamics and equilibrium computation to label spreading and residual feature propagation.

4.3 Model Properties Analysis

The convergence properties of the proposed model have been established in Theorem 1. We now proceed to analyze additional descriptive metrics of the model. In particular, the *price of anarchy* (PoA) quantifies the efficiency loss caused by decentralized decision-making, and is widely used to

measure the gap between the equilibrium outcome and the socially optimal solution. The following theorem provides an upper bound on the PoA of our model.

Theorem 2 (PoA upper bound for the normalized FJ model)**.** *Within the above defined framework, let the individual cost be defined as in* (8)*, and let the total social cost be* $C(\mathbf{z}) = \sum_{i=1}^{n} c_i(z_i)$*. Denote by* $\mathbf{z}^{\mathrm{NE}}$ *the Nash equilibrium and by* $\mathbf{z}^{\mathrm{OPT}}$ *the minimizer of* $C(\mathbf{z})$*. Then the price of anarchy satisfies*

$$\mathrm{PoA} \triangleq \frac{C(\mathbf{z}^{\mathrm{NE}})}{C(\mathbf{z}^{\mathrm{OPT}})} \leq \frac{9}{8}. \tag{9}$$

The degree-based change of variables used in the proof maps the normalized game to the standard undirected opinion game with positive node weights. Because this transformation preserves individual and social costs, the normalized formulation retains the same worst-case $9/8$ efficiency guarantee.

Another property of interest is the *conservation law of conflict* [12] in the vanilla FJ model (1):

$$\mathcal{C}_{\mathrm{int}} + \mathcal{C}_{\mathrm{ext}} + \mathcal{C} = \mathbf{s}^{\top}\mathbf{s}, \tag{10}$$

where $\mathcal{C}_{\mathrm{int}} = \sum_{i=1}^{n}(z_i - s_i)^2$ denotes the *internal conflict*, $\mathcal{C}_{\mathrm{ext}} = 2\sum_{\{i,j\}\in E} w_{ij}(z_i - z_j)^2$ denotes the *external conflict*, and $\mathcal{C} = \sum_{i=1}^{n} z_i^2$ represents the *controversy*. Clearly, the original conservation identity no longer holds after introducing the new parameters. Nevertheless, by incorporating appropriate weights, we can establish an analogous balanced relation under the new model, as stated in the following theorem.

Theorem 3 (Generalized normalized conflict identity)**.** *Let* $G = (V, E, \mathbf{W})$ *be a connected undirected weighted graph, and let* $\mathbf{z}^*$ *be the equilibrium of the normalized FJ model* $\mathbf{z}^* = \mathbf{D}_\alpha\mathbf{s} + (\mathbf{I} - \mathbf{D}_\alpha)\mathbf{S}\mathbf{z}^*$*, where* $\mathbf{D}_\alpha = \mathrm{diag}(\alpha_1, \ldots, \alpha_n)$ *with* $\alpha_i \in (0, 1)$*, and* $\mathbf{S} = \mathbf{D}^{-1/2}\mathbf{W}\mathbf{D}^{-1/2}$*. Define* $\bar{\mathcal{C}}_{\mathrm{int}} \triangleq \sum_{i\in V}\alpha_i(z_i^* - s_i)^2, \bar{\mathcal{C}}_\alpha \triangleq \sum_{i\in V}\alpha_i(z_i^*)^2$*, and*

$$\widetilde{\mathcal{C}}_{\mathrm{ext}} \triangleq 2\sum_{\{i,j\}\in E} w_{ij}\left((1-\alpha_i)\frac{z_i^*}{\sqrt{d_i}} - (1-\alpha_j)\frac{z_j^*}{\sqrt{d_j}}\right)\left(\frac{z_i^*}{\sqrt{d_i}} - \frac{z_j^*}{\sqrt{d_j}}\right).$$

Then the following exact identity holds: $\bar{\mathcal{C}}_{\mathrm{int}} + \widetilde{\mathcal{C}}_{\mathrm{ext}} + \bar{\mathcal{C}}_\alpha = \mathbf{s}^{\top}\mathbf{D}_\alpha\mathbf{s}$.

Moreover, if $\alpha_i \equiv \alpha$ *for all* $i \in V$*, then* $\widetilde{\mathcal{C}}_{\mathrm{ext}} = 2(1-\alpha)\sum_{\{i,j\}\in E} w_{ij}\left(\frac{z_i^*}{\sqrt{d_i}} - \frac{z_j^*}{\sqrt{d_j}}\right)^2$.

From Theorem 3, we observe that although the conservation law differs from its vanilla counterpart, the conflicts in the proposed model still maintain a delicate balance. This further validates the soundness and internal consistency of our model design.

5 Efficient Algorithms for Opinions Computation

5.1 Standard Synchronous Iterative Method

Within the framework of the FJ model, a variety of measures have been extensively studied, including polarization [67, 34, 55], disagreement [35], internal conflict and controversy [12], as well as co-exposure [19, 52]. Among these, the equilibrium opinion vector [50, 65] is arguably the most fundamental, as it underpins the computation and analysis of many of the aforementioned measures. A simple and direct way to compute the equilibrium opinion vector is via matrix inversion according to (6), which we hereafter refer to as Algorithm EXACT. This approach incurs a time complexity of $O(n^3)$ thus suffers from severe scalability limitations, making it impractical for large-scale networks and typically restricting its applicability to graphs with at most tens of thousands of nodes.

As an alternative approach for approximate computation, we note that Theorem 1 establishes that (5), whether interpreted as an opinion evolution process or as a fixed-point iteration, converges to the target vector $\mathbf{z}^*$ as the number of iterations approaches infinity. This observation naturally motivates an iterative method for approximating the equilibrium opinions. However, directly applying (5) as a fixed-point iteration leads to a nontrivial difficulty. In particular, the error analysis relies on the spectral radius of $(\mathbf{I} - \mathbf{D}_\alpha)\mathbf{S}$, whose non-symmetry complicates its connection to the ℓ_2-norm and necessitates an additional equivalent similarity transformation. To address these issues, we propose a

carefully designed similarity transformation that yields an equivalent fixed-point iteration with the same solution. The resulting formulation preserves the convergence rate while enabling a significantly sharper error analysis under the ℓ_2-norm.

We begin with (5) and observe that it is equivalent to the following transformed iteration:

$$(\mathbf{I}-\mathbf{D}_\alpha)^{-\frac{1}{2}}\mathbf{z}^{(t+1)} = (\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}\mathbf{S}(\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}(\mathbf{I}-\mathbf{D}_\alpha)^{-\frac{1}{2}}\mathbf{z}^{(t)} + (\mathbf{I}-\mathbf{D}_\alpha)^{-\frac{1}{2}}\mathbf{D}_\alpha\mathbf{s}.$$

Introducing the change of variables $\mathbf{y}^{(t)} = (\mathbf{I}-\mathbf{D}_\alpha)^{-\frac{1}{2}}\mathbf{z}^{(t)}$, we can rewrite the above iteration in the simplified form

$$\mathbf{y}^{(t+1)} = \mathbf{H}\mathbf{y}^{(t)} + \mathbf{c}, \tag{11}$$

where $\mathbf{H} = (\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}\mathbf{S}(\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}, \mathbf{c} = (\mathbf{I}-\mathbf{D}_\alpha)^{-\frac{1}{2}}\mathbf{D}_\alpha\mathbf{s}$. It is straightforward to verify that (11) and (5) share identical asymptotic convergence rates [44], since their iteration matrices are similar. Moreover, as $t \to \infty$, the sequence $\{\mathbf{y}^{(t)}\}$ converges to $\lim_{t\to\infty}\mathbf{y}^{(t)} \triangleq \mathbf{y}^* = (\mathbf{I}-\mathbf{D}_\alpha)^{-\frac{1}{2}}\mathbf{z}^*$.

Algorithm 1: ITERATIVE

Input: Graph $\mathcal{G}$, innate opinion vector $\mathbf{s}$, resistance parameters $\mathbf{D}_\alpha$, error tolerance ϵ
Output: Approximate equilibrium opinion vector $\hat{\mathbf{z}}$

1 $\mathbf{H} \leftarrow (\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}\mathbf{D}^{-\frac{1}{2}}\mathbf{W}\mathbf{D}^{-\frac{1}{2}}(\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}$, $\mathbf{c} \leftarrow (\mathbf{I}-\mathbf{D}_\alpha)^{-\frac{1}{2}}\mathbf{D}_\alpha\mathbf{s}$;
2 $\mathbf{y} \leftarrow \mathbf{c}$, $\mathbf{y}' \leftarrow \mathbf{0}$, $\alpha_{\text{res}} \leftarrow \min_{i\in\{1,2,\ldots,n\}}\mathbf{D}_\alpha[i,i]$;
3 **while** $\frac{1}{\alpha_{\text{res}}}\|\mathbf{y}-\mathbf{y}'\|_2 \geq \epsilon$ **do**
4 $\quad\mathbf{y}' \leftarrow \mathbf{y}$, $\mathbf{y} \leftarrow \mathbf{H}\mathbf{y} + \mathbf{c}$;
5 **return** $\hat{\mathbf{z}} = (\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}\mathbf{y}$

The pseudocode of the iterative approximation method based on (11) is presented in Algorithm 1, which we hereafter refer to as ITERATIVE. The error guarantee and computational complexity of Algorithm 1 are established in Theorem 4.

Theorem 4. *After at most* $t_{\max} = 1 + \max\left\{0, \left\lceil \frac{\log\left(\frac{\alpha_{\max}\sqrt{n}}{\alpha_{\min}\epsilon}\right)}{-\log(1-\alpha_{\min})} \right\rceil\right\}$ *iterations, Algorithm 1 returns an estimate* $\hat{\mathbf{z}}$ *such that* $\|\mathbf{z}^* - \hat{\mathbf{z}}\|_2 \leq \epsilon$. *The overall time complexity is bounded by* $O(mt_{\max})$.

It is worth noting that Theorem 4 provides an upper bound on the iteration complexity based on a relaxed analysis (see Appendix). In practice, since line 3 of Algorithm 1 performs real-time error monitoring, the actual number of iterations is typically significantly smaller than $t_{\max}$.

5.2 Conjugate Gradient Acceleration

In the above subsection, we derived an equivalent iterative formulation (11) that admits an explicit error guarantee. By denoting $\mathbf{M} = \mathbf{I} - \mathbf{H}$, we observe that $\mathbf{M}$ is symmetric and positive definite, which provides the necessary condition for applying the Conjugate Gradient (CG) method [24] to solve the corresponding linear system. The Conjugate Gradient method is a classical deterministic algorithm for solving symmetric positive-definite linear systems. Its key idea is to iteratively reduce the error by projecting it onto a sequence of mutually conjugate directions, ensuring efficient progress toward the solution. Throughout the iterations, the residual vectors remain mutually orthogonal, and in exact arithmetic, CG recovers the exact solution in at most n steps [45]. From a geometric perspective, CG can be interpreted as performing conjugate descent on a quadratic objective within a sequence of Krylov subspaces [28]. Further details can be found in a well-organized introduction [46].

We adopt the CG method instead of other linear solver here for two main reasons. First, it provides error bounds in the ℓ_2-norm or energy norm, which align well with our problem and introduce only mild relaxation, allowing for accurate error control. Second, the convergence rate of CG critically depends on the condition number of the system matrix $\mathbf{M}$. While preconditioning is typically required to improve conditioning, our construction of $\mathbf{H}$ in the previous subsection implicitly performs a Jacobi-like preconditioning, resulting in a well-conditioned $\mathbf{M}$. Consequently, directly applying CG to $\mathbf{M}$ yields fast convergence in practice.

Regarding error control, as in the standard iterative method, a conventional strategy is to derive an iteration upper bound from the spectral radius of the iteration matrix and the initial error, which

guarantees a prescribed accuracy. However, this strategy may lead to unnecessary over-computation. Instead, we exploit an invariant identity maintained throughout the CG iterations and terminate the algorithm using a runtime error criterion. Based on the above, the pseudocode of the CG method is presented in Algorithm 2, and its error guarantee and complexity analysis are summarized in Theorem 5. As further verified in the experimental section, the runtime error criterion often eliminates more than 50% of unnecessary iterations compared with using a fixed maximum iteration number.

Algorithm 2: CONJUGATE GRADIENT (CG)

Input: Graph $\mathcal{G}$, innate opinion vector $\mathbf{s}$, resistance parameters $\mathbf{D}_\alpha$, error tolerance ϵ
Output: Approximate equilibrium opinion vector $\hat{\mathbf{z}}$

$\mathbf{M} \leftarrow \mathbf{I} - (\mathbf{I} - \mathbf{D}_\alpha)^{\frac{1}{2}} \mathbf{D}^{-\frac{1}{2}} \mathbf{W} \mathbf{D}^{-\frac{1}{2}} (\mathbf{I} - \mathbf{D}_\alpha)^{\frac{1}{2}}$, $\mathbf{c} \leftarrow (\mathbf{I} - \mathbf{D}_\alpha)^{-\frac{1}{2}} \mathbf{D}_\alpha \mathbf{s}$;
$\mathbf{y}^{(0)} \leftarrow \mathbf{0}$, $\mathbf{r}^{(0)} \leftarrow \mathbf{c}$, $\mathbf{d}^{(0)} \leftarrow \mathbf{r}^{(0)}$, $t \leftarrow 0$;
$\alpha_{\text{res}} \leftarrow \min_i \mathbf{D}_\alpha[i,i]$, $t'_{\max} \leftarrow$ the bound in Theorem 5 using the prescribed resistance bounds;
while $\|\mathbf{r}^{(t)}\|_2 \geq \alpha_{\text{res}} \epsilon$ ***and*** $t < t'_{\max}$ **do**
  $a^{(t)} \leftarrow \frac{\mathbf{r}^{(t)\top} \mathbf{r}^{(t)}}{\mathbf{d}^{(t)\top} \mathbf{M} \mathbf{d}^{(t)}}$;
  $\mathbf{y}^{(t+1)} \leftarrow \mathbf{y}^{(t)} + a^{(t)} \mathbf{d}^{(t)}$;
  $\mathbf{r}^{(t+1)} \leftarrow \mathbf{r}^{(t)} - a^{(t)} \mathbf{M} \mathbf{d}^{(t)}$;
  $b^{(t)} \leftarrow \frac{\mathbf{r}^{(t+1)\top} \mathbf{r}^{(t+1)}}{\mathbf{r}^{(t)\top} \mathbf{r}^{(t)}}$;
  $\mathbf{d}^{(t+1)} \leftarrow \mathbf{r}^{(t+1)} + b^{(t)} \mathbf{d}^{(t)}$, $t \leftarrow t + 1$;
return $\hat{\mathbf{z}} = (\mathbf{I} - \mathbf{D}_\alpha)^{\frac{1}{2}} \mathbf{y}^{(t)}$

Theorem 5. *After at most* $t'_{\max} = \max\left\{0, \left\lceil \frac{\ln\left(\frac{2}{\epsilon}\sqrt{\frac{\tau n}{(1-\alpha_{\max})}}\right)}{\ln\left(\frac{\sqrt{\tau}+1}{\sqrt{\tau}-1}\right)} \right\rceil\right\}$, *where* $\tau = \frac{2-\alpha_{\min}}{\alpha_{\min}}$. *iterations, Algorithm 2 returns an estimate* $\hat{\mathbf{z}}$ *such that* $\|\mathbf{z}^* - \hat{\mathbf{z}}\|_2 \leq \epsilon$. *The overall time complexity is bounded by* $O(m(1 + t'_{\max}))$.

6 Experiments

Experimental Setting. Our algorithms are implemented in the Julia language and executed in single-threaded mode. All experiments are conducted on a MacBook Pro (2024) equipped with Apple M4 Max silicon and 36GB of main memory. We use a total of 12 real-world social network graphs from Konect [30] and a special Reddit network [14] with real-world collected innate opinions in our experiments. During simulation, innate opinions are randomly generated within the interval $[0, 1]$, and resistance parameters are sampled from the interval $[0.01, 0.99]$ using three distributions: a uniform distribution, a power-law distribution with density $f(\alpha) = \frac{C}{\alpha^{1.5}}$, where $C = \frac{1}{2}\left(\alpha_{\min}^{-0.5} - \alpha_{\max}^{-0.5}\right)^{-1}$ is the normalization constant, and a bimodal distribution with two modes concentrated near $\alpha_{\min}$ and $\alpha_{\max}$, respectively.

6.1 Model Analysis

We first conduct a comparative analysis of different models on the real-world Reddit network dataset with empirically collected opinion data [14]. In total, we consider four models: the widely used vanilla FJ model [9] (2), the conflict-normalized FJ model [11, 2] (4), the simple edge-weight normalized FJ model (7), and the fully normalized FJ model (5). The four variants distinguish resistance parameterization from interaction normalization. Our formulation decouples the innate-opinion coefficient from node degree and normalizes pairwise influence using both endpoints' degrees. At fixed resistance parameters, the conflict-balanced and fully normalized formulations differ in one-endpoint versus two-endpoint normalization. We consider the opinion dynamics to have reached equilibrium when $\|\mathbf{z}^{(t+1)} - \mathbf{z}^{(t)}\|_2 \leq 10^{-8}$. Based on the equilibrium state, we report several descriptive system-level metrics, including the average opinion $\bar{z} = \frac{1}{n}\mathbf{1}^\top \mathbf{z}$ and the number of iterations required for convergence. In addition, we summarize four structural measures—internal conflict $\mathcal{I} = \sum_{i \in V} (z_i - s_i)^2$, disagreement $\mathcal{D} = \mathbf{z}^\top \mathbf{L} \mathbf{z}$, normalized disagreement $\mathcal{D}_{\text{norm}} = \mathbf{z}^\top \mathbf{L}_{\text{sym}} \mathbf{z}$, and polarization $\mathcal{P} = \sum_{i \in V} \tilde{z}_i^2 = \tilde{\mathbf{z}}^\top \tilde{\mathbf{z}}$, *where* $\tilde{\mathbf{z}} = \mathbf{z} - \bar{z}\mathbf{1}$ denotes the centered

opinion vector. Table 1 reports the results under all three resistance distributions. The vanilla and edge-normalized baselines remain unchanged across distributions because they do not depend on the sampled resistance parameters.

Table 1: Comparison of FJ variants on Reddit across resistance distributions.

Distribution	Model	Iterations	$\bar{z}^*$	$\mathcal{I}$	$\mathcal{D}$	$\mathcal{D}_{\mathrm{norm}}$	$\mathcal{P}$
Uniform	Vanilla FJ	342	0.498151	0.854972	0.048329	10.243075	0.010951
	Conflict balanced FJ	29	0.497981	0.324469	9.170614	10.562975	0.311717
	Edge normalized FJ	23	0.480555	2.793483	50.958648	2.554223	2.635907
	Full normalized FJ	31	0.480655	3.561183	65.311631	3.497478	3.313792
Power-law	Vanilla FJ	342	0.498151	0.854972	0.048329	10.243075	0.010951
	Conflict balanced FJ	132	0.499125	0.804513	0.621693	10.291051	0.030026
	Edge normalized FJ	23	0.480555	2.793483	50.958648	2.554223	2.635907
	Full normalized FJ	140	0.469861	8.936374	133.282484	0.277772	7.679263
Bimodal	Vanilla FJ	342	0.498151	0.854972	0.048329	10.243075	0.010951
	Conflict balanced FJ	20	0.499594	0.425040	12.794347	10.987897	0.463810
	Edge normalized FJ	23	0.480555	2.793483	50.958648	2.554223	2.635907
	Full normalized FJ	21	0.479778	4.942487	84.875954	5.074575	5.029134

Based on Table 1, several observations can be drawn. We first compare the vanilla FJ model with the edge-weight normalized variant. Both models do not incorporate resistance parameters and differ only in the use of normalization. We observe that normalization breaks the preservation of the average opinion, while significantly accelerating convergence, as evidenced by the reduced number of iterations. At the same time, both internal conflict and polarization increase, indicating that normalization, while reweighting pairwise interactions, may amplify deviations from intrinsic opinions and lead to more dispersed outcomes.

A key observation concerns the behavior of disagreement. Since normalization alters the definition of pairwise interactions, directly comparing the standard disagreement $\mathcal{D}$ across models is not fully appropriate. To provide a consistent comparison, we also report the normalized disagreement $\mathcal{D}_{\mathrm{norm}}$. Notably, the two measures exhibit opposite trends across models, reflecting that they capture different aspects of opinion divergence under distinct interaction mechanisms. Across all three reported distributions, the fully normalized model has lower normalized disagreement but higher polarization than the conflict-balanced model. The distinction is particularly clear under power-law resistance, where low normalized disagreement coexists with substantial global opinion dispersion. These results highlight that agreement in degree-normalized coordinates and dispersion around the global mean describe different aspects of the equilibrium. Across all experimental settings, both Theorem 2 and Theorem 3 are consistently validated by the empirical results.

With respect to the resistance parameter α, both the conflict-balanced and fully normalized models require more iterations under the power-law distribution than under the uniform distribution, whereas the bimodal distribution requires fewer. The power-law profile places greater weight on small resistance parameters, making individuals more susceptible to external opinions, consistent with the slower convergence observed in this setting. Under bimodal resistance, the fully normalized model exhibits both higher normalized disagreement and higher polarization than under uniform resistance, showing that faster convergence can coexist with greater equilibrium opinion dispersion. We are not the first to introduce resistance parameters into the FJ model; a comprehensive analysis of how resistance parameters influence opinion dynamics can be found in [2, 1].

6.2 Opinion Computation

In this subsection, we evaluate both the efficiency and effectiveness of the proposed algorithms. As baselines, we consider EXACT and ITERATIVE, introduced in Section 5.1. Although ITERATIVE is treated as a baseline, our contribution also includes an improved version of it. Specifically, instead of relying on a pre-specified maximum number of iterations $t_{\max}$ [1], we incorporate a runtime error check to adaptively terminate the iteration once the desired accuracy is reached, thereby effectively reducing redundant computation. We clarify why several other approaches for computing equilibrium opinions are not adopted as baselines. Random sampling–based methods, such as Monte Carlo techniques [20] and random spanning forest sampling [50], rely on connections to absorbing random walks or the fundamental matrix via spanning forests [6]. Adapting these methods to our normalized

formulation requires reformulating the system and translating their error guarantees to the original opinion coordinates. Another line of work employs SDDM solvers [48, 65]. Our formulation enables CG to operate directly on the symmetric positive-definite system $\mathbf{My} = \mathbf{c}$, using sparse matrix–vector products and explicit error control in the original opinion coordinates.

Table 2: Running time (seconds) and approximation error under uniform resistance parameters.

Network	n	m	Exact	Iterative				CG			
				Time	it	$it_{\max}$	Error (10^{-8})	Time	it	$it_{\max}$	Error (10^{-8})
Hamsterster households	1,576	4,032	0.028	0.016	1708	2658	0.942	0.001	122	196	0.100
Hamsterster friendships	1,858	12,534	0.043	0.002	122	2666	0.044	0.001	31	196	0.006
Hamsterster full	2,426	16,631	0.078	0.008	303	2679	0.980	0.001	53	197	0.162
Facebook (NIPS)	2,888	12,534	0.126	0.001	60	2688	0.010	0.000	20	198	0.002
Brightkite	58,228	214,078	-	0.262	467	2837	-	0.105	70	209	-
Catster households	105,138	494,858	-	2.039	1885	2867	-	0.286	162	211	-
Douban	154,908	327,162	-	0.141	79	2886	-	0.046	33	212	-
Catster friends	204,473	5,448,197	-	12.861	1905	2900	-	1.276	164	213	-
Youtube friendships	1,134,890	2,987,624	-	2.253	210	2985	-	0.704	49	219	-
Flixster	2,523,386	7,918,801	-	3.291	126	3025	-	1.404	42	222	-
Orkut	3,072,441	117,184,899	-	13.381	55	3035	-	7.750	28	223	-
LiveJournal	10,690,276	112,307,385	-	578.873	2115	3097	-	56.773	177	227	-

Table 3: Running time (seconds) and approximation error under power-law resistance parameters.

Network	n	m	Exact	Iterative				CG			
				Time	it	$it_{\max}$	Error (10^{-8})	Time	it	$it_{\max}$	Error (10^{-8})
Hamsterster households	1,576	4,032	0.027	0.016	1901	2658	0.954	0.001	160	196	0.169
Hamsterster friendships	1,858	12,534	0.041	0.028	1788	2666	0.959	0.002	121	196	0.070
Hamsterster full	2,426	16,631	0.074	0.034	1642	2679	0.660	0.003	146	197	0.127
Facebook (NIPS)	2,888	12,534	0.123	0.009	363	2688	0.121	0.000	22	198	0.001
Brightkite	58,228	214,078	-	0.832	1779	2837	-	0.168	164	209	-
Catster households	105,138	494,858	-	2.353	2088	2867	-	0.306	178	211	-
Douban	154,908	327,162	-	0.396	320	2886	-	0.094	67	212	-
Catster friends	204,473	5,448,197	-	14.379	2108	2900	-	1.420	180	213	-
Youtube friendships	1,134,890	2,987,624	-	6.936	663	2985	-	1.348	98	219	-
Flixster	2,523,386	7,918,801	-	12.228	492	3025	-	2.653	85	222	-
Orkut	3,072,441	117,184,899	-	62.198	264	3035	-	14.585	54	223	-
LiveJournal	10,690,276	112,307,385	-	639.370	2305	3097	-	62.676	193	227	-

We report the running time and approximation error under uniform and power-law resistance in Tables 2 and 3, respectively, with randomly generated initial opinions and $\epsilon = 10^{-8}$; results under bimodal resistance are provided in Appendix F. The theoretical iteration limits in Tables 2, 3, and 4 use the common bounds $\alpha_{\min} = 0.01$ and $\alpha_{\max} = 0.99$. It is worth noting that, for ITERATIVE, there exists a substantial gap between the empirical iteration count and the theoretical bound. This observation confirms that our proposed runtime error-checking mechanism significantly improves practical efficiency by avoiding a large number of redundant iterations. When $n > 5 \times 10^4$, the running time of EXACT exceeds two hours. In contrast, CG achieves a substantial speedup, with the improvement becoming more pronounced on larger networks. On small-scale networks, we measure the approximation error using the ℓ_2 norm. The results show that both iterative methods achieve the desired high-precision solutions. Under power-law resistance, CG still requires fewer iterations than ITERATIVE on every tested network and achieves the prescribed accuracy on the smaller networks.

7 Conclusion

In this work, we proposed a normalized variant of the classical Friedkin–Johnsen model that incorporates degree information from both endpoints to achieve a more balanced and principled characterization of interpersonal influence. The resulting formulation not only preserves the essential properties of FJ dynamics, but also induces a new equilibrium behavior with improved structural interpretability. We established the convergence of the proposed model, analyzed its intrinsic properties from both spectral and game-theoretic perspectives, and developed efficient algorithms for computing equilibrium opinions and related descriptive metrics. Extensive experiments on large-scale networks demonstrated the effectiveness and scalability of our approach.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Nos. 62372112).

References

[1] Rediet Abebe, T-H Hubert Chan, Jon Kleinberg, Zhibin Liang, David Parkes, Mauro Sozio, and Charalampos E Tsourakakis. Opinion dynamics optimization by varying susceptibility to persuasion via non-convex local search. *ACM Transactions on Knowledge Discovery from Data*, 16(2):1–34, 2021.

[2] Rediet Abebe, Jon Kleinberg, David Parkes, and Charalampos E Tsourakakis. Opinion dynamics with varying susceptibility to persuasion. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1089–1098, 2018.

[3] Daron Acemoglu, Giacomo Como, Fabio Fagnani, and Asuman Ozdaglar. Opinion fluctuations and disagreement in social networks. *Mathematics of Operations Research*, 38(1):1–27, 2013.

[4] Eugene Agichtein, Carlos Castillo, Debora Donato, Aristides Gionis, and Gilad Mishne. Finding high-quality content in social media. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*, pages 183–194, 2008.

[5] Brian DO Anderson and Mengbin Ye. Recent advances in the modelling and analysis of opinion dynamics on influence networks. *International Journal of Automation and Computing*, 16(2):129–149, 2019.

[6] Luca Avena and Alexandre Gaudillière. Two applications of random spanning forests. *Journal of Theoretical Probability*, 31(4):1975–2004, 2018.

[7] Alain Barrat, Marc Barthélemy, Romualdo Pastor-Satorras, and Alessandro Vespignani. The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences*, 101(11):3747–3752, 2004.

[8] Carmela Bernardo, Lingfei Wang, Francesco Vasca, Yiguang Hong, Guodong Shi, and Claudio Altafini. Achieving consensus in multilateral international negotiations: The case study of the 2015 Paris Agreement on climate change. *Science Advances*, 7(51):eabg8068, 2021.

[9] David Bindel, Jon Kleinberg, and Sigal Oren. How bad is forming your own opinion? *Games and Economic Behavior*, 92:248–265, 2015.

[10] T-H. Hubert Chan, Zhibin Liang, and Mauro Sozio. Revisiting opinion dynamics with varying susceptibility to persuasion via non-convex local search. In *Proceedings of the 2019 World Wide Web Conference*, pages 173–183, 2019.

[11] Xi Chen, Jefrey Lijffijt, and Tijl De Bie. The normalized friedkin-johnsen model (a work-in-progress report). In *ECML PKDD 2018-PhD Forum*, 2018.

[12] Xi Chen, Jefrey Lijffijt, and Tijl De Bie. Quantifying and minimizing risk of conflict in social networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1197–1205, 2018.

[13] Federico Cinus, Atsushi Miyauchi, Yuko Kuroki, and Francesco Bonchi. Minimizing polarization and disagreement in the Friedkin–Johnsen model with unknown innate opinions. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pages 7365–7373, 2025.

[14] Abir De, Sourangshu Bhattacharya, Parantapa Bhattacharya, Niloy Ganguly, and Soumen Chakrabarti. Learning a linear influence model from transient opinion dynamics. In *Proceedings of the 23rd ACM Int. Conf. Info. Knowl. Manage.*, pages 401–410, 2014.

[15] Morris H DeGroot. Reaching a consensus. *Journal of the American Statistical association*, 69(345):118–121, 1974.

[16] Yucheng Dong, Min Zhan, Gang Kou, Zhaogang Ding, and Haiming Liang. A survey on the fusion process in opinion dynamics. *Information Fusion*, 43:57–65, 2018.

[17] Noah E Friedkin and Eugene C Johnsen. Social influence and opinions. *Journal of Mathematical Sociology*, 15(3-4):193–206, 1990.

[18] Noah E Friedkin and Eugene C Johnsen. *Social influence network theory: A sociological examination of small group dynamics*, volume 33. Cambridge University Press, 2011.

[19] Kiran Garimella, Aristides Gionis, Nikos Parotsidis, and Nikolaj Tatti. Balancing information exposure in social networks. *Advances in neural information processing systems*, 30, 2017.

[20] Aristides Gionis, Evimaria Terzi, and Panayiotis Tsaparas. Opinion maximization in social networks. In *Proceedings of the 2013 SIAM International Conference on Data Mining*, pages 387–395, 2013.

[21] Patrick Groeber, Jan Lorenz, and Frank Schweitzer. Dissonance minimization as a microfoundation of social influence in models of opinion formation. *The Journal of Mathematical Sociology*, 38(3):147–174, 2014.

[22] Guang He, Wenbing Zhang, Jing Liu, and Haoyue Ruan. Opinion dynamics with the increasing peer pressure and prejudice on the signed graph. *Nonlinear Dynamics*, 99:3421–3433, 2020.

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

[24] Magnus R Hestenes and Eduard Stiefel. Methods of conjugate gradients for solving linear systems. *Journal of research of the National Bureau of Standards*, 49(6):409–436, 1952.

[25] Richard A Holley and Thomas M Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability*, pages 643–663, 1975.

[26] Petter Holme, Beom Jun Kim, Chang No Yoon, and Seung Kee Han. Attack vulnerability of complex networks. *Physical Review E*, 65(5):056109, 2002.

[27] Peng Jia, Anahita MirTabatabaei, Noah E Friedkin, and Francesco Bullo. Opinion dynamics and the evolution of social power in influence networks. *SIAM Review*, 57(3):367–397, 2015.

[28] Carl T Kelley. *Iterative Methods for Linear and Nonlinear Equations*. SIAM, 1995.

[29] Dilip Krishnan, Raanan Fattal, and Richard Szeliski. Efficient preconditioning of laplacian matrices for computer graphics. *ACM Transactions on Graphics (TOG)*, 32(4):1–15, 2013.

[30] Jérôme Kunegis. Konect: The Koblenz Network Collection. In *Proceedings of the 22nd International World Wide Web Conference*, pages 1343–1350, 2013.

[31] Heidi Ledford. How Facebook, Twitter and other data troves are revolutionizing social science. *Nature*, 582(7812):328–330, 2020.

[32] Hao Li, Hao Jiang, Yuke Zheng, Hao Sun, and Wenying Gong. Unigo: A unified graph neural network for modeling opinion dynamics on graphs. In *Proceedings of the ACM on Web Conference 2025*, pages 530–540, 2025.

[33] Patrick Jackson Macdonald, Eivind Almaas, and A.-L. Barabási. Minimum spanning trees of weighted scale-free networks. *EPL (Europhysics Letters)*, 72(2):308–314, 2005.

[34] Antonis Matakos, Evimaria Terzi, and Panayiotis Tsaparas. Measuring and moderating opinion polarization in social networks. *Data Mining and Knowledge Discovery*, 31:1480–1505, 2017.

[35] Cameron Musco, Christopher Musco, and Charalampos E Tsourakakis. Minimizing polarization and disagreement in social networks. In *Proceedings of the 2018 World Wide Web Conference*, pages 369–378, 2018.

[36] Stefan Neumann, Yinhao Dong, and Pan Peng. Sublinear-time opinion estimation in the friedkin–johnsen model. In *Proceedings of the ACM Web Conference 2024*, pages 2563–2571, 2024.

[37] Daniele Notarmuzi, Claudio Castellano, Alessandro Flammini, Dario Mazzilli, and Filippo Radicchi. Universality, criticality and complexity of information propagation in social media. *Nature Communications*, 13:1308, 2022.

[38] Charlotte Out, Sijing Tu, Stefan Neumann, and Ahad N Zehmakan. The impact of external sources on the Friedkin–Johnsen model. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 1815–1824, 2024.

[39] Sergey E Parsegov, Anton V Proskurnikov, Roberto Tempo, and Noah E Friedkin. Novel multidimensional models of opinion dynamics in social networks. *IEEE Transactions on Automatic Control*, 62(5):2270–2285, 2017.

[40] Anton V Proskurnikov and Roberto Tempo. A tutorial on modeling and analysis of dynamic social networks. Part I. *Annual Reviews in Control*, 43:65–79, 2017.

[41] Anton V Proskurnikov, Roberto Tempo, Ming Cao, and Noah E Friedkin. Opinion evolution in time-varying social influence networks with prejudiced agents. *IFAC-PapersOnLine*, 50(1):11896–11901, 2017.

[42] Inzamam Rahaman and Patrick Hosein. Extending Degroot opinion formation for signed graphs and minimizing polarization. In *Complex Networks & Their Applications IX: Volume 2, Proceedings of the Ninth International Conference on Complex Networks and Their Applications COMPLEX NETWORKS 2020*, pages 298–309, 2021.

[43] Roberta Raineri, Giacomo Como, and Fabio Fagnani. Optimal selection of the most informative nodes in opinion dynamics on networks. *IFAC-PapersOnLine*, 56(2):4192–4197, 2023.

[44] Yousef Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, 2003.

[45] Timothy Sauer. *Numerical Analysis*. Addison-Wesley Publishing Company, 2011.

[46] Jonathan Richard Shewchuk. An introduction to the conjugate gradient method without the agonizing pain. Technical report, School of Computer Science, Carnegie Mellon University, 1994.

[47] Mohammad Shirzadi, Ali Safarpoor Dehkordi, and Ahad N Zehmakan. Adaptive initial residual connections for gnns with theoretical guarantees. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 25392–25400, 2026.

[48] Daniel A Spielman and Shang-Hua Teng. Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM Journal on Matrix Analysis and Applications*, 35(3):835–885, 2014.

[49] Samuel Stern and Giacomo Livan. The impact of noise and topology on opinion dynamics in social networks. *Royal Society Open Science*, 8(4):201943, 2021.

[50] Haoxin Sun and Zhongzhi Zhang. Opinion optimization in directed social networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 4623–4632, 2023.

[51] Panayiotis Tsaparas, Alexandros Ntoulas, and Evimaria Terzi. Selecting a comprehensive set of reviews. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 168–176, 2011.

[52] Sijing Tu, Cigdem Aslay, and Aristides Gionis. Co-exposure maximization in online social networks. *Advances in Neural Information Processing Systems*, 33:3232–3243, 2020.

[53] Gengyu Wang, Runze Zhang, and Zhongzhi Zhang. Efficient algorithms for relevant quantities of friedkin-johnsen opinion dynamics model. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 2915–2926, 2025.

[54] Gengyu Wang, Runze Zhang, and Zhongzhi Zhang. Opinion maximization in social networks by modifying internal opinions. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen, editors, *Advances in Neural Information Processing Systems*, volume 38, pages 77209–77238, 2025.

[55] Yanbang Wang and Jon Kleinberg. On the relationship between relevance and conflict in online social link recommendations. *Advances in Neural Information Processing Systems*, 36:36708–36725, 2023.

[56] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *International conference on machine learning*, pages 6861–6871. Pmlr, 2019.

[57] Pinghua Xu, Wenbin Hu, Jia Wu, and Weiwei Liu. Opinion maximization in social trust networks. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence*, pages 1251–1257, 2020.

[58] Wanyue Xu, Qi Bao, and Zhongzhi Zhang. Fast evaluation for relevant quantities of opinion dynamics. In *Proceedings of The Web Conference*, pages 2037–2045, 2021.

[59] Wanyue Xu, Yubo Sun, Mingzhe Zhu, Zuobai Zhang, and Zhongzhi Zhang. Behavior and sublinear algorithm for opinion disagreement on noisy social networks. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–15, 2026.

[60] Zuobai Zhang, Wanyue Xu, Zhongzhi Zhang, and Guanrong Chen. Opinion dynamics incorporating higher-order interactions. In *Proceedings of the 20th IEEE International Conference on Data Mining*, pages 1430–1435, 2020.

[61] Zuobai Zhang, Wanyue Xu, Zhongzhi Zhang, and Guanrong Chen. Opinion dynamics in social networks incorporating higher-order interactions. *Data Mining and Knowledge Discovery*, 38(6):4001–4023, 2024.

[62] Haoran Zheng, Renchi Yang, Yubo Zhou, and Jianliang Xu. Rethinking message passing neural networks with diffusion distance-guided stress majorization. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pages 2031–2041, 2026.

[63] Dengyong Zhou, Olivier Bousquet, Thomas Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. *Advances in neural information processing systems*, 16, 2003.

[64] Dengyong Zhou and Bernhard Schölkopf. Regularization on discrete spaces. In *Joint Pattern Recognition Symposium*, pages 361–368. Springer, 2005.

[65] Xiaotian Zhou, Haoxin Sun, Wanyue Xu, Wei Li, and Zhongzhi Zhang. Friedkin-Johnsen model for opinion dynamics on signed graphs. *IEEE Transactions on Knowledge and Data Engineering*, 36(12):8313–8327, 2024.

[66] Xiaotian Zhou, Liwang Zhu, Wei Li, and Zhongzhi Zhang. A sublinear time algorithm for opinion optimization in directed social networks via edge recommendation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3593–3602, 2023.

[67] Liwang Zhu, Qi Bao, and Zhongzhi Zhang. Minimizing polarization and disagreement in social networks via link recommendation. *Advances in Neural Information Processing Systems*, 34:2072–2084, 2021.

[68] Liwang Zhu and Zhongzhi Zhang. Opinion maximization in social networks via link recommendation. *Theoretical Computer Science*, page 115090, 2025.

[69] Liwang Zhu, Xiaotian Zhou, Jiahe Tian, Wei Li, and Zhongzhi Zhang. Defending against malicious influence control in online leader-follower social networks. *IEEE Transactions on Information Forensics and Security*, 19:4809–4819, 2024.

A Proof of Theorem 1

Proof. Let $\mathbf{M} \triangleq (\mathbf{I} - \mathbf{D}_\alpha)\mathbf{S}$, then the dynamics in (5) can be written as $\mathbf{z}(t+1) = \mathbf{D}_\alpha \mathbf{s} + \mathbf{M}\mathbf{z}(t)$. To prove convergence, it therefore suffices to show that $\rho(\mathbf{M}) < 1$ [44].

To this end, define $\mathbf{B} \triangleq (\mathbf{I} - \mathbf{D}_\alpha)^{1/2}$. Since $\alpha_i \in (0,1)$ for all $i \in V$, $\mathbf{B}$ is a well-defined diagonal matrix satisfying $\|\mathbf{B}\|_2^2 = \max_{i\in V}(1-\alpha_i) < 1$. Consider $\mathbf{H} \triangleq \mathbf{BSB}$. Because $\mathbf{S}$ is symmetric, $\mathbf{H}$ is also symmetric. Moreover, since $\mathbf{B}^{-1}\mathbf{M}\mathbf{B} = \mathbf{B}^{-1}\mathbf{B}^2\mathbf{S}\mathbf{B} = \mathbf{B}\mathbf{S}\mathbf{B} = \mathbf{H}$, and hence $\mathbf{M}$ and $\mathbf{H}$ are similar and thus share the same eigenvalues. Next, since $\mathbf{S} = \mathbf{D}^{-1/2}\mathbf{W}\mathbf{D}^{-1/2}$ is the normalized adjacency matrix of an undirected graph, all its eigenvalues lie in $[-1,1]$, and thus $\|\mathbf{S}\|_2 \le 1$. Therefore,

$$\rho(\mathbf{M}) = \rho(\mathbf{H}) \le \|\mathbf{H}\|_2 = \|\mathbf{BSB}\|_2 \le \|\mathbf{B}\|_2^2 \|\mathbf{S}\|_2 \le \max_{i\in V}(1-\alpha_i) < 1.$$

It follows from standard linear iteration theory [28] that $\mathbf{M}^t \to \mathbf{0}$ as $t \to \infty$, and hence the sequence $\{\mathbf{z}(t)\}_{t\ge 0}$ converges to the unique fixed point

$$\mathbf{z}^* = \big(\mathbf{I} - (\mathbf{I} - \mathbf{D}_\alpha)\mathbf{S}\big)^{-1}\mathbf{D}_\alpha \mathbf{s}.$$

□

B Proof of Theorem 2

Proof. Define

$$x_i \triangleq \frac{z_i}{\sqrt{d_i}}, \qquad q_i \triangleq \frac{s_i}{\sqrt{d_i}}. \tag{12}$$

Then $z_i = \sqrt{d_i}\, x_i$ and $s_i = \sqrt{d_i}\, q_i$, so that

$$(z_i - s_i)^2 = d_i (x_i - q_i)^2.$$

Substituting this identity into the individual cost (8), we obtain

$$c_i(\mathbf{x}) = \beta_i (x_i - q_i)^2 + \sum_{j\in N_i} w_{ij}(x_i - x_j)^2, \tag{13}$$

where

$$\beta_i \triangleq \frac{\alpha_i d_i}{1-\alpha_i} > 0. \tag{14}$$

This bijective change of variables preserves each player's cost, so best responses and Nash equilibria correspond exactly between the two formulations. Moreover, it is precisely the standard undirected Friedkin–Johnsen game with arbitrary positive node weights β_i.

Let the corresponding social cost be

$$C(\mathbf{x}) = \sum_{i=1}^n \beta_i (x_i - q_i)^2 + 2 \sum_{\{i,j\}\in E} w_{ij}(x_i - x_j)^2. \tag{15}$$

By Claim 3.5 of [9], for undirected graphs with arbitrary node weights, the price of anarchy is at most $9/8$, i.e.,

$$C(\mathbf{x}^{\mathrm{NE}}) \le \frac{9}{8}\, C(\mathbf{x}^{\mathrm{OPT}}).$$

Finally, since the mapping $z_i = \sqrt{d_i}\, x_i$ is bijective, the Nash equilibrium and the social optimum correspond exactly between the $\mathbf{z}$- and $\mathbf{x}$-variables, and the social cost is preserved under this change of variables. Therefore,

$$\frac{C(\mathbf{z}^{\mathrm{NE}})}{C(\mathbf{z}^{\mathrm{OPT}})} \le \frac{9}{8}.$$

This completes the proof. □

C Proof of Theorem 3

Proof. For notational simplicity, write $\mathbf{z} = \mathbf{z}^*$, and define

$$\mathbf{A} \triangleq \mathbf{D}_\alpha, \qquad \mathbf{B} \triangleq \mathbf{I} - \mathbf{D}_\alpha, \qquad \mathbf{H} \triangleq \mathbf{I} - \mathbf{S}.$$

Then $\mathbf{A}$ and $\mathbf{B}$ are diagonal matrices, and since $\mathbf{W}$ is symmetric, the matrix $\mathbf{S} = \mathbf{D}^{-1/2}\mathbf{W}\mathbf{D}^{-1/2}$ is symmetric. Therefore, $\mathbf{H}$ is also symmetric.

The equilibrium equation

$$\mathbf{z} = \mathbf{A}\mathbf{s} + \mathbf{B}\mathbf{S}\mathbf{z}$$

can be rearranged as

$$\mathbf{A}(\mathbf{s} - \mathbf{z}) = \mathbf{B}(\mathbf{I} - \mathbf{S})\mathbf{z} = \mathbf{B}\mathbf{H}\mathbf{z}.$$

That is,

$$\mathbf{A}(\mathbf{s} - \mathbf{z}) = \mathbf{B}\mathbf{H}\mathbf{z}. \tag{16}$$

Next, expanding the quadratic form $\mathbf{s}^\top \mathbf{A}\mathbf{s}$ around $\mathbf{z}$, we obtain

$$\mathbf{s}^\top \mathbf{A}\mathbf{s} = (\mathbf{z} + (\mathbf{s} - \mathbf{z}))^\top \mathbf{A}(\mathbf{z} + (\mathbf{s} - \mathbf{z})).$$

Hence,

$$\mathbf{s}^\top \mathbf{A}\mathbf{s} = \mathbf{z}^\top \mathbf{A}\mathbf{z} + 2\mathbf{z}^\top \mathbf{A}(\mathbf{s} - \mathbf{z}) + (\mathbf{s} - \mathbf{z})^\top \mathbf{A}(\mathbf{s} - \mathbf{z}).$$

Using (16), this becomes

$$\mathbf{s}^\top \mathbf{A}\mathbf{s} = \mathbf{z}^\top \mathbf{A}\mathbf{z} + 2\mathbf{z}^\top \mathbf{B}\mathbf{H}\mathbf{z} + (\mathbf{s} - \mathbf{z})^\top \mathbf{A}(\mathbf{s} - \mathbf{z}). \tag{17}$$

It remains to rewrite the middle term in edge form. Let

$$\mathbf{x} \triangleq \mathbf{D}^{-1/2}\mathbf{z}, \qquad \mathbf{L} \triangleq \mathbf{D} - \mathbf{W}.$$

Since

$$\mathbf{H} = \mathbf{I} - \mathbf{S} = \mathbf{D}^{-1/2}\mathbf{L}\mathbf{D}^{-1/2},$$

and $\mathbf{B}$ commutes with $\mathbf{D}^{-1/2}$ because both are diagonal, we have

$$\mathbf{z}^\top \mathbf{B}\mathbf{H}\mathbf{z} = \mathbf{x}^\top \mathbf{B}\mathbf{L}\mathbf{x}.$$

Therefore,

$$2\mathbf{z}^\top \mathbf{B}\mathbf{H}\mathbf{z} = 2\mathbf{x}^\top \mathbf{B}\mathbf{L}\mathbf{x}. \tag{18}$$

Now use the standard Laplacian identity for undirected weighted graphs: for any vectors $\mathbf{u}, \mathbf{v}$,

$$\mathbf{u}^\top \mathbf{L}\mathbf{v} = \sum_{\{i,j\} \in E} w_{ij}(u_i - u_j)(v_i - v_j).$$

Applying this identity with $\mathbf{u} = \mathbf{B}\mathbf{x}$ and $\mathbf{v} = \mathbf{x}$ yields

$$\mathbf{x}^\top \mathbf{B}\mathbf{L}\mathbf{x} = (\mathbf{B}\mathbf{x})^\top \mathbf{L}\mathbf{x} = \sum_{\{i,j\} \in E} w_{ij}\big((\mathbf{B}\mathbf{x})_i - (\mathbf{B}\mathbf{x})_j\big)(x_i - x_j).$$

Since $(\mathbf{B}\mathbf{x})_i = (1 - \alpha_i)x_i$, it follows that

$$\mathbf{x}^\top \mathbf{B}\mathbf{L}\mathbf{x} = \sum_{\{i,j\} \in E} w_{ij}\left((1 - \alpha_i)x_i - (1 - \alpha_j)x_j\right)(x_i - x_j).$$

Substituting $x_i = z_i/\sqrt{d_i}$ into (18), we obtain

$$2\mathbf{z}^\top \mathbf{B}\mathbf{H}\mathbf{z} = 2\sum_{\{i,j\} \in E} w_{ij}\left((1 - \alpha_i)\frac{z_i}{\sqrt{d_i}} - (1 - \alpha_j)\frac{z_j}{\sqrt{d_j}}\right)\left(\frac{z_i}{\sqrt{d_i}} - \frac{z_j}{\sqrt{d_j}}\right) = \widetilde{\mathcal{C}}_{\text{ext}}.$$

Substituting this into (17), and recalling the definitions of $\bar{\mathcal{C}}_{\text{int}}$ and $\bar{\mathcal{C}}_\alpha$, we conclude that

$$\bar{\mathcal{C}}_{\text{int}} + \widetilde{\mathcal{C}}_{\text{ext}} + \bar{\mathcal{C}}_\alpha = \mathbf{s}^\top \mathbf{D}_\alpha \mathbf{s}.$$

Finally, if $\alpha_i \equiv \alpha$ for all i, then

$$(1 - \alpha_i)\frac{z_i}{\sqrt{d_i}} - (1 - \alpha_j)\frac{z_j}{\sqrt{d_j}} = (1 - \alpha)\left(\frac{z_i}{\sqrt{d_i}} - \frac{z_j}{\sqrt{d_j}}\right),$$

and hence

$$\widetilde{\mathcal{C}}_{\text{ext}} = 2(1 - \alpha)\sum_{\{i,j\} \in E} w_{ij}\left(\frac{z_i}{\sqrt{d_i}} - \frac{z_j}{\sqrt{d_j}}\right)^2.$$

Thus the identity reduces to the homogeneous conservation law. The proof is complete. □

D Proof of Theorem 4

Let $\alpha_{\text{res}} = \min_i \alpha_i \geq \alpha_{\min}$. Write $\mathbf{y}^{(0)} = \mathbf{c}$ for the initialization in Algorithm 1 and set $\mathbf{y}^{(-1)} = \mathbf{0}$. From (11), define

$$\Delta\mathbf{y}^{(t+1)} \triangleq \mathbf{y}^{(t+1)} - \mathbf{y}^{(t)}.$$

Then

$$\Delta\mathbf{y}^{(t+1)} = \mathbf{H}\Delta\mathbf{y}^{(t)}.$$

Hence, for any $\ell \geq 0$,

$$\Delta\mathbf{y}^{(t+1+\ell)} = \mathbf{H}^{\ell}\Delta\mathbf{y}^{(t+1)}.$$

Since $\mathbf{y}^{(t)} \to \mathbf{y}^*$ as $t \to \infty$, we have

$$\mathbf{y}^* - \mathbf{y}^{(t)} = \sum_{i=t+1}^{\infty} \Delta\mathbf{y}^{(i)} = \sum_{\ell=0}^{\infty} \mathbf{H}^{\ell}\Delta\mathbf{y}^{(t+1)} = (\mathbf{I}-\mathbf{H})^{-1}\Delta\mathbf{y}^{(t+1)},$$

where the last equality follows from $\rho(\mathbf{H}) < 1$.

Recalling that

$$\mathbf{z}^* = (\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}\mathbf{y}^*,$$

we obtain

$$\mathbf{z}^* - (\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}\mathbf{y}^{(t)} = (\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}(\mathbf{I}-\mathbf{H})^{-1}\Delta\mathbf{y}^{(t+1)}.$$

Taking the ℓ_2-norm on both sides yields

$$\left\|\mathbf{z}^* - (\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}\mathbf{y}^{(t)}\right\|_2 \leq \|(\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}\|_2 \|(\mathbf{I}-\mathbf{H})^{-1}\|_2 \|\Delta\mathbf{y}^{(t+1)}\|_2.$$

Since $\mathbf{H}$ is symmetric, we have

$$\|\mathbf{H}\|_2 = \rho(\mathbf{H}) \leq 1 - \alpha_{\text{res}}.$$

Therefore,

$$\|(\mathbf{I}-\mathbf{H})^{-1}\|_2 \leq \frac{1}{1-\|\mathbf{H}\|_2} \leq \frac{1}{\alpha_{\text{res}}}.$$

Moreover,

$$\|(\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}\|_2 = \sqrt{\|\mathbf{I}-\mathbf{D}_\alpha\|_2} \leq \sqrt{1-\alpha_{\text{res}}}.$$

Combining the above bounds, we obtain

$$\left\|\mathbf{z}^* - (\mathbf{I}-\mathbf{D}_\alpha)^{\frac{1}{2}}\mathbf{y}^{(t)}\right\|_2 \leq \frac{\sqrt{1-\alpha_{\text{res}}}}{\alpha_{\text{res}}}\|\Delta\mathbf{y}^{(t+1)}\|_2 \leq \frac{1}{\alpha_{\text{res}}}\|\Delta\mathbf{y}^{(t+1)}\|_2.$$

Since $\Delta\mathbf{y}^{(t+1)} = \mathbf{H}\Delta\mathbf{y}^{(t)}$ and $\|\mathbf{H}\|_2 \leq 1 - \alpha_{\text{res}} < 1$, this also gives

$$\left\|\mathbf{z}^* - (\mathbf{I}-\mathbf{D}_\alpha)^{1/2}\mathbf{y}^{(t)}\right\|_2 \leq \frac{1}{\alpha_{\text{res}}}\|\Delta\mathbf{y}^{(t)}\|_2.$$

Thus the difference tested in Algorithm 1 guarantees the required accuracy, including at initialization.

To bound the number of iterations, write $q = 1 - \alpha_{\min}$. After k loop iterations, the tested difference is

$$\Delta\mathbf{y}^{(k)} = \mathbf{H}^k\mathbf{c}, \qquad \Delta\mathbf{y}^{(0)} = \mathbf{c}.$$

Moreover,

$$\mathbf{Hc} = (\mathbf{I}-\mathbf{D}_\alpha)^{1/2}\mathbf{S}\mathbf{D}_\alpha\mathbf{s}, \qquad \|\mathbf{Hc}\|_2 \leq \sqrt{q}\,\alpha_{\max}\sqrt{n},$$

where we used $\|\mathbf{S}\|_2 \leq 1$ and $\|\mathbf{s}\|_2 \leq \sqrt{n}$. Hence, for $k \geq 1$,

$$\frac{\|\Delta\mathbf{y}^{(k)}\|_2}{\alpha_{\text{res}}} \leq \frac{q^{k-1}\|\mathbf{Hc}\|_2}{\alpha_{\min}} \leq q^{k-1/2}\frac{\alpha_{\max}\sqrt{n}}{\alpha_{\min}}.$$

For

$$k = t_{\max} = 1 + \max\left\{0, \left\lceil \frac{\log\left(\frac{\alpha_{\max}\sqrt{n}}{\alpha_{\min}\epsilon}\right)}{-\log q} \right\rceil\right\},$$

the last bound is at most $\sqrt{q}\,\epsilon < \epsilon$. Thus Algorithm 1 stops after at most $t_{\max}$ iterations. Each iteration costs $O(m)$, yielding the stated $O(mt_{\max})$ complexity.

E Proof of Theorem 5

Write $\alpha_{\mathrm{res}} = \min_i \alpha_i \geq \alpha_{\min}$ as in Algorithm 2. During the iteration process, we note that:

$$\mathbf{M}\mathbf{y}^{(t+1)} + \mathbf{r}^{(t+1)} = \mathbf{M}(\mathbf{y}^{(t)} + a^{(t)}\mathbf{d}^{(t)}) + \mathbf{r}^{(t)} - a^{(t)}\mathbf{M}\mathbf{d}^{(t)} = \mathbf{M}\mathbf{y}^{(t)} + \mathbf{r}^{(t)}.$$

By induction, we can deduce that $\mathbf{r}^{(t)} = \mathbf{c} - \mathbf{M}\mathbf{y}^{(t)} = \mathbf{M}(\mathbf{y}^* - \mathbf{y}^{(t)})$, and

$$\|(\mathbf{I} - \mathbf{D}_\alpha)^{\frac{1}{2}}\mathbf{y}^{(t)} - \mathbf{z}^*\|_2 = \|(\mathbf{I} - \mathbf{D}_\alpha)^{\frac{1}{2}}\mathbf{M}^{-1}\mathbf{r}^{(t)}\|_2 \leq \|\mathbf{M}^{-1}\|_2\|\mathbf{r}^{(t)}\|_2 \leq \frac{1}{\alpha_{\mathrm{res}}}\|\mathbf{r}^{(t)}\|_2.$$

The last inequality holds because $\|(\mathbf{I} - \mathbf{D}_\alpha)^{\frac{1}{2}}\|_2 \leq 1$ and $\|\mathbf{M}^{-1}\|_2 \leq \frac{1}{\alpha_{\mathrm{res}}}$. Thus the residual criterion $\|\mathbf{r}^{(t)}\|_2 < \alpha_{\mathrm{res}}\epsilon$ guarantees the required accuracy.

To obtain a time complexity bound, denote $\mathbf{e}^{(t)} = \mathbf{y}^{(t)} - \mathbf{y}^*$, and then we have $\|\mathbf{e}^{(t)}\|_{\mathbf{M}} \leq 2(\frac{\sqrt{\kappa_{\mathbf{M}}}-1}{\sqrt{\kappa_{\mathbf{M}}}+1})^t\|\mathbf{e}^{(0)}\|_{\mathbf{M}}$, proved in [46], Section 9.2. To bound the initial energy, observe that

$$\mathbf{M} - \mathbf{D}_\alpha = (\mathbf{I} - \mathbf{D}_\alpha)^{1/2}\mathbf{L}_{\mathrm{sym}}(\mathbf{I} - \mathbf{D}_\alpha)^{1/2} \succeq \mathbf{0}.$$

Since $\mathbf{D}_\alpha \succ \mathbf{0}$, this implies $\mathbf{M}^{-1} \preceq \mathbf{D}_\alpha^{-1}$. Using $\mathbf{y}^{(0)} = \mathbf{0}$ and $\mathbf{M}\mathbf{y}^* = \mathbf{c}$, we obtain

$$\begin{aligned}\|\mathbf{e}^{(0)}\|_{\mathbf{M}}^2 &= \mathbf{c}^\top\mathbf{M}^{-1}\mathbf{c} \leq \mathbf{c}^\top\mathbf{D}_\alpha^{-1}\mathbf{c} \\ &= \mathbf{s}^\top\mathbf{D}_\alpha(\mathbf{I} - \mathbf{D}_\alpha)^{-1}\mathbf{s} \leq \frac{\alpha_{\max}n}{1 - \alpha_{\max}} \leq \frac{(2 - \alpha_{\min})n}{1 - \alpha_{\max}}.\end{aligned}$$

Similarly, we have

$$\|\mathbf{e}^{(t)}\|_{\mathbf{M}} = \|\mathbf{y}^{(t)} - \mathbf{y}^*\|_{\mathbf{M}} \geq \sqrt{\lambda_{\mathbf{M}\min}}\|\mathbf{y}^{(t)} - \mathbf{y}^*\|_2 \geq \sqrt{\alpha_{\min}}\|\mathbf{y}^{(t)} - \mathbf{y}^*\|_2.$$

By $\|\mathbf{e}^{(t)}\|_{\mathbf{M}} \leq 2(\frac{\sqrt{\kappa_{\mathbf{M}}}-1}{\sqrt{\kappa_{\mathbf{M}}}+1})^t\|\mathbf{e}^{(0)}\|_{\mathbf{M}}$, we can derive:

$$\|(\mathbf{I} - \mathbf{D}_\alpha)^{\frac{1}{2}}\mathbf{y}^{(t)} - \mathbf{z}^*\|_2 \leq \|\mathbf{y}^{(t)} - \mathbf{y}^*\|_2 \leq 2(\frac{\sqrt{\kappa_{\mathbf{M}}} - 1}{\sqrt{\kappa_{\mathbf{M}}} + 1})^t\sqrt{\frac{(2 - \alpha_{\min})n}{(1 - \alpha_{\max})\alpha_{\min}}} = f(t).$$

Note that $\kappa_{\mathbf{M}} = \frac{\lambda_{\mathbf{M}\max}}{\lambda_{\mathbf{M}\min}} \leq \frac{2-\alpha_{\min}}{\alpha_{\min}} = \tau$. Replacing $\kappa_{\mathbf{M}}$ by τ in the preceding bound shows that the error is at most ϵ after

$$t'_{\max} = \max\left\{0, \left\lceil \frac{\ln\left(\frac{2}{\epsilon}\sqrt{\frac{\tau n}{1-\alpha_{\max}}}\right)}{\ln\left(\frac{\sqrt{\tau}+1}{\sqrt{\tau}-1}\right)} \right\rceil\right\}$$

iterations. If Algorithm 2 stops earlier, its residual criterion guarantees the same accuracy; otherwise, it stops at $t = t'_{\max}$ with the fixed-iteration guarantee. Including preprocessing, the total time complexity is $O(m(1+t'_{\max}))$, since each loop is dominated by a sparse matrix–vector multiplication.

F Ablation Study under Bimodal Resistance on the Larger Networks

Table 4: Running time (seconds) and approximation error under bimodal resistance parameters.

Network	n	m	Exact	Iterative				CG			
				Time	it	$it_{\max}$	Error (10^{-8})	Time	it	$it_{\max}$	Error (10^{-8})
Hamsterster households	1,576	4,032	0.027	0.017	2032	2658	0.976	0.001	137	196	0.251
Hamsterster friendships	1,858	12,534	0.037	0.027	1806	2666	0.969	0.001	42	196	0.006
Hamsterster full	2,426	16,631	0.089	0.036	1893	2679	0.870	0.003	113	197	0.011
Facebook (NIPS)	2,888	12,534	0.122	0.001	73	2688	0.014	0.000	16	198	0.006
Brightkite	58,228	214,078	-	0.879	1960	2837	-	0.155	151	209	-
Catster households	105,138	494,858	-	2.437	2232	2867	-	0.294	172	211	-
Douban	154,908	327,162	-	0.178	114	2886	-	0.061	42	212	-
Catster friends	204,473	5,448,197	-	15.325	2253	2900	-	1.361	173	213	-
Youtube friendships	1,134,890	2,987,624	-	3.866	358	2985	-	1.076	75	219	-
Flixster	2,523,386	7,918,801	-	5.483	213	3025	-	1.843	58	222	-
Orkut	3,072,441	117,184,899	-	36.043	150	3035	-	10.976	38	223	-
LiveJournal	10,690,276	112,307,385	-	680.355	2456	3097	-	58.346	184	227	-

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The claims made match the theoretical and experimental results in the paper.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: This paper models social networks as general undirected weighted graphs. While more complex structures—such as hypergraphs and multidimensional networks—lead to different discussions and algorithmic designs, they are beyond the scope of current work and are left for future research.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All theorems and lemmas are properly proved or cited.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Reproducibility is supported by the explicit model equations, Algorithms 1–2, the proofs in Appendices A–E, and the data sources, parameter settings, stopping criteria, and computing environment reported in Section 6.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: A public code package is not included with the paper. The algorithms and experimental setup are described in the paper, and the sources of the network datasets are cited.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: All experimental details have been fully presented in the Experimental Section.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: The proposed algorithm is deterministic; therefore, running it multiple times under the same conditions yields identical results. We provide ablation studies under various parameter settings, to demonstrate the performance and robustness of the method.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The PC data is reported, and its specifications can be found on the manufacturer's official website.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research presented in this paper conforms to the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: This work presents a foundational improvement in modeling opinion dynamics and the corresponding algorithm design. The primary impact is enhancing computational and analytical efficiency for researchers in related fields. We believe the work poses minimal risk of direct negative societal impact.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The code in this work is mathematical and computational in nature, and it utilizes publicly available datasets. Consequently, there is minimal to no risk of misuse associated with its release.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All external assets used in this work are appropriately cited, and their respective licenses and terms of use are clearly acknowledged.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.

- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: No new asset is released.

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.