

UNIVERSITY OF CALIFORNIA
SANTA BARBARA

**Analysis, Modeling, and Control of
Dynamic Processes in Networks**

A dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy
in
Computer Science

by

Victor Amelkin

Committee in charge:

Professor Ambuj K. Singh, Chair
Professor Francesco Bullo
Professor John R. Gilbert
Professor Xifeng Yan

June 2018

The Dissertation of Victor Amelkin is approved:

Professor Francesco Bullo

Professor John R. Gilbert

Professor Xifeng Yan

Professor Ambuj K. Singh, Committee Chair

May 2018

Analysis, Modeling, and Control of
Dynamic Processes in Networks

Copyright © 2018

by

Victor Amelkin

To my wife and kids

Acknowledgements

Throughout the course of my PhD, I have had a pleasure of interacting with, learning from, and being influenced and supported by a multitude of people whose contribution to my growing as a scientist and a person I would like to acknowledge here.

Foremost, I express gratitude to my research advisor Ambuj K. Singh for being a great mentor, academic role model, and supporter of my work. I have enormously benefited from Ambuj's sharing his technical knowledge, as well as his perspective on conducting scientific research and successfully functioning as a part of the academic community. Over the passed 5+ years of working with Ambuj and due to his influence, I have shaped my research vision, and grown as a scientist and a person. I am also grateful to Amr El Abbadi and Divy Agrawal, who, in response to my proposal to work on theoretical aspects of distributed systems in early 2013, suggested that I chat to Ambuj; luckily, that chat back then eventually led to my doctorate.

I am grateful to Francesco Bullo, from collaboration with whom I have greatly benefited, and from the meetings with whom over the past three years I have learned a lot about the realm of dynamical systems and control theory. Due to Francesco, I have better understood the link between Computer Science and Engineering, and integrated both into my research vision. I am also grateful to John R. Gilbert, over a few months of work with whom I had explored the connection between graph algorithms and linear algebra, which—mixed together with my learning modern matrix algebra from Esmond G. Ng and Shiv Chandrasekaran—paved the way for me towards studying dynamic processes in networks. I attribute my interest in studying networked social systems to my acquaintance and interaction with Noah E. Friedkin. In addition to his contributing ideas to some of my published work, he has demonstrated to me the beauty of mathematical sociology, shared his wisdom, and served as an inspiration for my research. I also extend my gratitude to Xifeng Yan for his serving on my doctoral committee, as well as his academic advice and encouragement.

It was a privilege to collaborate with Petko Bogdanov, Young Ji Kim, Thomas W. Malone, and Omid Askarisichani, with whom we co-authored works on theoretical and empirical network process analysis. My gratitude further extends to all the members of the Network Science of Teams MURI for their collegiality. I am especially grateful to Brian Uzzi, who has piqued my interest in studying financial networks, served as an inspiration and a role model of an academic leader, and whose advice helped me define my future career path.

I enjoyed meeting with and tangibly benefited from the advice of affiliates and visitors of UCSB's Center for Control, Dynamical Systems and Computation. In particular, I am grateful to Roberto Tempo (1956-2017), Fabio Fagnani, Jason Marden, and Rakesh Vohra for their advice. I have also enjoyed meeting great network scientists over the past months, who provided me with a new outlook upon the field of network science, and influenced my career choices.

I am delighted to have spent the last several years with my fellow graduate students and postdocs from UCSB's Dynamo Lab, Francesco's research group, Network Science Laboratory, CCDC, and Combinatorial Scientific Computing Lab.

Finally, my warmest and deepest appreciation goes to my wife Maria, my son Fedor, and other members of my family. Maria has been a great partner and supporter of my academic endeavors over the years, and Fedor has been consistently filling my life with joy as of December, 2013. Without them, this thesis would not be possible.

My research has been largely supported by the US Army Research Lab and the US Army Research Office through the Network Science of Teams MURI grant W911NF-15-1-0577, and to a lesser extent by the same agencies through the Institute for Collaborative Biotechnologies grant W911NF-09-0001, as well as through the Querying and Mining of Dynamic Graphs grant IIS-1219254 from NSF.

Curriculum Vitæ

Victor Amelkin

Education

2011-2018 PhD in Computer Science
University of California, Santa Barbara

2006-2008 MSc in Applied Mathematics and Computer Science
Tula State University, Tula, Russia

2002-2006 BSc in Applied Mathematics and Computer Science
Tula State University, Tula, Russia

Experience

06/2013-06/2018 *Graduate Student Researcher*, UC Santa Barbara, CA

06/2016-09/2016 *Research Scientist Intern*, Amazon.com, Seattle, WA

07/2014-09/2014 *Research Intern*, Technicolor Labs, Los Altos, CA

07/2012-09/2012 *Research Assistant*, Lawrence Berkeley National Lab, Berkeley, CA

10/2011-06/2012 *Graduate Student Researcher*, UC Santa Barbara, CA

10/2006-03/2010 *Software Engineer*, DevExpress, Tula, Russia

Professional Service

Served as a reviewer for a number of computer science and engineering conferences and journals, including:

	KDD	SDM	WWW	ICDM	WSDM	AAAI	TKDE	TNSE	TAC	Automatica	ACC	CDC
2018	•		•				•		•	•	•	
2017	•	•	•			•	•	•				•
2016	•	•		•	•	•						
2015	•						•					
2014							•					
2013				•								

Teaching

10/2012-09/2013 *Teaching Assistant*, UC Santa Barbara, CA

09/2007-06/2008 *Teaching Assistant*, Tula State University, Tula, Russia

Awards

2011-2012	\$8,000 Merit Fellowship, Computer Science Department, UCSB
2011	\$1,500 Opportunity Grant, US Department of State
2008	Best Paper, Conference on Current Problems of Mathematics, Mechanics, Computer Science, Tula, Russia
2002	Tula Regional Olympiad in Applied Mathematics for high-school students, 2 nd prize

Publications

Refereed Journals and Conferences:

- [1] **Amelkin, V.**, Bogdanov, P., Singh, A.K., “*A Distance Measure for the Analysis of Polar Opinion Dynamics in Social Networks (Extended Paper)*” // In submission to ACM TKDD (May, 2018).
- [2] **Amelkin, V.**, Singh, A.K., “*Disabling External Influence in Social Networks via Edge Recommendation*” // In submission to IEEE ICDD (May, 2018).
- [3] **Amelkin, V.**, Bullo, F., Singh, A.K., “*Polar Opinion Dynamics in Social Networks*” // IEEE Transactions on Automatic Control (61) 11, 2017.
- [4] **Amelkin, V.**, Askarisichani, O., Kim Y.J., Malone T.W., Singh, A.K., “*Dynamics of Collective Performance in Collaboration Networks*” // In submission to PLOS ONE (May, 2017).
- [5] **Amelkin, V.**, Bogdanov, P., Singh, A.K., “*A Distance Measure for the Analysis of Polar Opinion Dynamics in Social Networks*” // In Proc. of IEEE ICDE, May, 2017.

Working Papers:

- [6] **Amelkin, V.**, et al., “*Dynamics of Opinions and Social Power in Complex Social Networks*” // 2018.
- [7] **Amelkin, V.**, et al., “*A Model for Performance-Driven Collaboration Among Stock Traders*” // 2018.

Non-Refereed (Pre-prints, Workshops, Reports, Presentations, Theses):

- [8] **Amelkin, V.**, Singh, A.K., “*Disabling External Influence in Social Networks via Edge Recommendation (Extended Paper)*” // arXiv:1709.08139 [cs.SI], Sept., 2017.
- [9] **Amelkin, V.**, Bullo, F., Singh, A.K., “*Polar Opinion Dynamics in Social Networks*” // 32nd Southern California Control Workshop, Caltech, April, 2017.
- [10] **Amelkin, V.**, Askarisichani, O., Singh, A.K., Kim Y.J., Malone T.W., “*Dynamics of Collective Performance in Collaboration Networks*” // INSNA Sunbelt, Newport Beach, April, 2016.
- [11] **Amelkin, V.**, Singh, A.K., Bogdanov, P., “*A Distance Measure for the Analysis of Polar Opinion Dynamics in Social Networks (Extended Paper)*” // arXiv:1510.05058 [cs.SI], 2015.
- [12] **Amelkin, V.**, Ng, E.G., “*Parallel Communication Analysis for Sparse Cholesky Factorization Algorithms*” // Technical Report, Lawrence Berkeley National Laboratory, 2012.
- [13] **Amelkin, V.**, “*Fejér Problem for Polynomials on a Sphere*” // MSc Thesis, Tula State University, 2008.

Abstract

Analysis, Modeling, and Control of Dynamic Processes in Networks

by

Victor Amelkin

Dynamic network processes have surrounded people for millennia. Information spread through social networks, alliance formation in financial and organizational networks, heat diffusion through material networks, and distributed synchronization in robotic networks are just a few examples. Network processes are studied along three dimensions: analysis of network processes through the data produced by them; designing complex plausible, yet, tractable mathematical models for network processes; and designing control mechanisms that would guide network processes towards desirable evolution patterns. This thesis advances the frontier of knowledge about network processes along each of these three dimensions, emphasizing applications to social networks.

The first part of the thesis is dedicated to the design of a method for model-driven analysis of a polar opinion formation process in social networks. The core of the method is a distance measure quantifying the likelihood of a social network's transitioning between different states with respect to a chosen opinion dynamics model characterizing expected evolution of the network's state. I describe how to design such a distance measure relying upon the classical transportation problem, compute it in linear time, and use it in applications.

In the second part of the thesis, I focus on designing a model for polar opinion formation in social networks, and define a class of non-linear models that capture the dependence of the users' opinion formation behavior upon the opinions themselves. The obtained models are connected to socio-psychological theories, and their behavior is theoretically analyzed employing tools from non-smooth analysis and a generalization of LaSalle Invariance Principle.

The third part of the thesis targets the problem of defense against social control. While the existing socio-psychological theories as well as influence maximization techniques expose the opinion formation process in social networks to external attacks, I propose an algorithm that nullifies the effect of such attacks by strategically recommending a small number of new edges to the network's users. The optimization problem underlying the algorithm is NP-hard, and I provide a pseudo-linear time heuristic—drawing upon the theory of Markov chains—that solves the problem approximately and performs well in experiments.

Published Content, Contributions, and Permissions

Some of the materials presented in this thesis have either been published by the thesis' author or are currently in submission. The author has made principal contributions to all stages of the conception and production of the published and submitted works mentioned below.

A fraction of the content of Chapter 2 has been published as

V. Amelkin, P. Bogdanov, A. K. Singh, *A Distance Measure for the Analysis of Polar Opinion Dynamics in Social Networks*, in Proceedings of IEEE International Conference on Data Engineering (ICDE), pp. 159–162, © IEEE, 2017, DOI: 10.1109/ICDE.2017.64

and a large part of that chapter's content is currently in submission. A major part of the content of Chapter 3 has been published as

V. Amelkin, F. Bullo, A. K. Singh, *Polar Opinion Dynamics in Social Networks*, *IEEE Transactions on Automatic Control*, 62 (11), © IEEE, 2017, DOI: 10.1109/TAC.2017.2694341

Finally, a large part of the content of Chapter 4 is currently in submission.

The IEEE does not require obtaining explicit permission for the reuse of the above mentioned content in this thesis. In reference to this IEEE copyrighted content, the IEEE does not endorse any of the University of California's products or services. Internal or personal use of this material is permitted. If interested in reprinting or republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please contact the IEEE.

Contents

Acknowledgements	v
Curriculum Vitæ	vii
Abstract	ix
Published Content, Contributions, and Permissions	xi
List of Figures	xv
List of Tables	xviii
1 Introduction	1
1.1 Static Network Science	2
1.2 Dynamic Network Science	4
1.3 Network Processes: Analysis, Modeling, Control	5
1.4 Overview of the Present Work	8
2 Model-Driven Analysis of Polar Opinion Dynamics in Social Networks	11
2.1 Introduction	12
2.2 Preliminaries	14
2.2.1 Network and Network States	14
2.2.2 User Opinions and Their Dynamics	15
2.2.3 Earth Mover’s Distance	17
2.3 Related Work	18
2.4 Social Network Distance	20
2.5 Generalized Earth Mover’s Distance	26
2.5.1 Why Do We Need a New Earth Mover’s Distance?	26
2.5.2 Generalized Earth Mover’s Distance	31
2.6 Efficient Computation of Social Network Distance	36
2.7 Experimental Results	42
2.7.1 Experimental Setup	42

2.7.2	Detecting Anomalous Network States	44
2.7.3	Predicting User Opinions	48
2.7.4	Sensitivity to Opinion Dynamics Models	52
2.7.5	Scalability of Social Network Distance	53
2.8	Limitations	54
2.9	Conclusion	55
2.10	Future Research	56
3	Non-Linear Models for Polar Opinion Dynamics in Social Networks	57
3.1	Introduction	58
3.2	Related Work	63
3.2.1	Analytic Models	63
3.2.2	Algorithmic Models	66
3.3	Models for Polar Opinion Dynamics	67
3.3.1	General Model	67
3.3.2	Specialized Models	68
3.4	Preliminaries from Non-smooth Analysis	70
3.5	Analysis of the Models	74
3.5.1	Well-posedness	75
3.5.2	Equilibrium Points	76
3.5.3	Convergence Analysis	81
3.6	Simulation Results	98
3.7	Discussion	98
4	Defense Against Social Control via Link Recommendation	105
4.1	Introduction	106
4.2	Preliminaries	110
4.3	Problem's Statement and Hardness	111
4.4	Related Work	117
4.4.1	Analytic Optimization of Network Topology	117
4.4.2	Combinatorial Optimization of Network Topology	119
4.4.3	Centrality Perturbation and Manipulation	121
4.5	Strategic Edge Addition to the Network	122
4.5.1	Selection of Candidate Edge Source Nodes	123
4.5.2	Eigencentality Under Single-Edge Perturbation	124
4.5.3	Asymptotic Consensus Value Under Single-Edge Perturbation	127
4.5.4	Efficient Computation of Candidate Edge Scores	128
4.5.5	Solving DIVER	132
4.5.6	Probabilistic Edge Addition	133
4.6	Experimental Results	133
4.6.1	Experimental Setup	134
4.6.2	Solving DIVER	136

4.6.3	DIVER’s Effectiveness and Efficiency	138
4.6.4	DIVER’s Robustness to Network Noise	140
4.7	Conclusion	140
4.8	Future Work	141
5	Conclusion	142
5.1	Open Problems in Network Process Analysis	143
5.2	Open Problems in Network Process Modeling	144
5.3	Open Problems in Network Process Control	146
5.4	Future of Network Process Research	147
	Bibliography	148

List of Figures

2.1	$d_{\mathcal{M}}(P, Q)$ accumulates likelihoods of all possible network state evolution paths.	21
2.2	Growth of the number of active users—holding non-neutral opinions with respect to topic “Obama”—in a sample of Twitter network over time.	25
2.3	Network states P and Q defined over the same network are extended with bank nodes, whose values are chosen, so that the total masses of the extended network states \tilde{P} and \tilde{Q} are equal. The ground distances $\tilde{D}_{bank,i} = \tilde{D}_{i,bank} = \gamma$ from and to the bank node are uniformly defined based on the largest ground distance between the initially present nodes.	27
2.4	Two qualitatively different ways to transport a unit of mass from extended network state $\tilde{P} = [P_1, \dots, P_n, k + \Delta]$ to extended network state $\tilde{Q} = [Q_1, \dots, Q_n, k]$, where $\sum P_i \leq \sum Q_j$. Dashed arrows represent the flow of mass. The bank node is attached to every node of each network state. $k = \sum P_i$, so that the total masses of two extended network states are equal.	29
2.5	A network state $[P_1, \dots, P_4, P^{(1)}, \dots, P^{(4)}]$ extended with <i>local bank nodes</i> $P^{(i)}$. The undirected edges to and from the bank nodes—displayed dashed—are weighted with ground distances $\gamma_1, \dots, \gamma_4$, respectively.	31
2.6	Three network states defined over the same two-cluster network.	32
2.7	Re-routing of mass flow to convert an optimal transportation plan \hat{f} into plan f^\dagger	38
2.8	Anomaly detection on synthetic data. $ V = 20k$, scale-free exponent $\gamma = -2.3$. A series of 40 network states is generated using $\mathbb{P}_{nbr} = 0.12$ and $\mathbb{P}_{ext} = 0.01$ for normal and $\mathbb{P}_{nbr} = 0.08$ and $\mathbb{P}_{ext} = 0.05$ for anomalous network states’ generation, respectively. The three simulated anomalies are displayed as solid vertical lines.	45
2.9	ROC curves comparing the quality of anomaly detection by different distance measures in a series of 300 network states over synthetic network with $ V = 30k$ and scale-free exponent $\gamma = -2.3$. The network states are generated using $\mathbb{P}_{nbr} = 0.08$ and $\mathbb{P}_{ext} = 0.001$ for normal and $\mathbb{P}_{nbr} = 0.07$ and $\mathbb{P}_{ext} = 0.011$ for anomalous instances.	46

2.10	Anomaly detection on Twitter data (May'08-Aug'11). The distance series are accompanied by the curve showing Google Trends' scaled interest in topic "Obama". Network states detected to be anomalous by at least one distance measure are displayed as solid vertical lines.	47
2.11	Distance measure-based user opinion prediction. Network states $G(t)$ reside in the network state space. The series of distances $\langle \dots, d_{-3}, d_{-2}, d_{-1} \rangle$ between past network states $\dots, G(-4), G(-3), G(-2), G(-1)$ adjacent in time is extrapolated to estimate the distance d^* to the true current network state $\widehat{G}(0)$ to be predicted. We, then, search for the assignment of opinions to target users in $G(0)$ to make the distance from $G(-1)$ to the obtained network state $G^*(0)$ as close to d^* as possible.	48
2.12	SND and ℓ_1 distances between network states of normal (ICC) and anomalous (random) transitions.	53
2.13	Time for computing SND when the number of users having different opinion is fixed at $n_\Delta = 1000$ and the total number of users n in the network grows up to 200k.	54
2.14	Time for computing SND using our method, with the network size fixed at $n = 20k$, and the number n_Δ of users who changed their opinions growing up to 10k.	54
3.1	Convergence behavior of the model with stubborn positives in two dimensions, as well as the partition of the state space. Several trajectories representative of the model's behavior are displayed as solid arrows.	87
3.2	Convergence behavior of the model with stubborn neutrals in two dimensions, as well as the partition of the state space. Several trajectories representative of the model's behavior are shown as solid arrows.	90
3.3	Convergence behavior of the model with stubborn extremists in three dimensions. A few representative trajectories are displayed as solid arrows; the solid diagonal correspond to the consensus equilibrium states of the system; the circles correspond to the disagreement equilibrium states; the corners are the states of disagreement where all agents are closed, while in the states in the interior of the cube's edges, some agents remain open (depending on the location of closed agents in the network, some edges may have no such internal points of equilibrium, like in the case of edge $[-1, -1, 1] - [-1, 1, 1]$).	93
3.4	Convergence of the model with stubborn extremists in $S_0(\epsilon)$	94
3.5	Convergence of the model with stubborn extremists in $S_-(\epsilon)$	94
3.6	Convergence of the model with stubborn extremists in $S_*(I_{\max}, I_{\min})$	95
3.7	Phase portrait: convergence to consensus of the model with stubborn positives in the absence of closed agents.	100
3.8	Phase portrait: convergence to consensus of the model with stubborn neutrals when closed agents are present.	101

3.9	Phase portrait: asymptotic behavior of the model with stubborn extremists. Strictly inside the state space, trajectories converge to a consensus. On the surface of the state space, when all closed agents have the same state α , trajectories converge to that consensus state $\alpha\mathbb{1}$. In the presence of multiple closed agents holding different opinions, the system converges to a state of disagreement (such states are displayed as solid circles).	102
4.1	DIVER’s definition by example. The adversary influences user opinions, $x \rightarrow \tilde{x}$, increasing the asymptotic consensus value $\langle \pi, x \rangle \rightarrow \langle \pi, \tilde{x} \rangle$. The network responds by adding edges, reducing the asymptotic consensus value $\langle \tilde{\pi}, \tilde{x} \rangle$, driving it back to original $\langle \pi, x \rangle$	108
4.2	Perturbation of a network with a single edge (r, c) with weight θ_{rc} . Existing edges are downweighted, maintaining row-stochasticity of the adjacency matrix.	112
4.3	Network W^{KC} for $n = 3$; absent edges C are displayed dashed. Node states $z \otimes \mathbb{1}_2$, used in the reduction, are displayed next to the nodes.	116
4.4	Dependence of the asymptotic consensus value’s reduction $f_\pi(r, c) = \langle \pi, \tilde{x} \rangle - \langle \tilde{\pi}, \tilde{x} \rangle$ after addition of edge (r, c) , $\theta_{rc} = \text{const}$ to a scale-free network ($n = 100, \gamma = -2.5$) upon the eigenvector centrality π_r of the edge’s source node.	123
4.5	Comparison of exact and approximate candidate edge scores f_π in a scale-free network ($n = 100, \gamma = -2.5$).	128
4.6	Root mean square error of MFPT estimates in a scale-free network ($n = 100, \gamma = -2.5$) using walks of different length.	130
4.7	Dependency of the length of a random walk—used for estimating MFPTs to and from top-centrality nodes—on the size and density of the scale-free network.	131
4.8	Solving DIVER over synthetic and real-world networks.	137
4.9	Quality of candidate edges added by DIVER($f_\pi \sim 20\%$) to SF(256, -2.5).	138
4.10	Dependence of DIVER($f_\pi \sim X\%$) performance on X over SF(256, -2.5).	139
4.11	Time for a single evaluation of our heuristic DIVER($f_\pi \sim \text{const}$), adding $k = 10$ edges.	139
4.12	Performance of DIVER($f_\pi \sim 100\%$) over SF(1024, -2.5) with exact and noisy edge weights.	140

List of Tables

2.1	Summary of notation used in the design of Social Network Distance	15
2.2	Means μ and standard deviations σ of user opinion prediction accuracies.	51
3.1	Simulation results for different models and initial states.	105
4.1	Summary of notation.	111
4.2	Networks used in experiments with DIVER.	135

Chapter 1

Introduction

In the world we live in, objects rarely exist and function in isolation, yet, they typically are connected to and interact with others. The fundamental mathematical concept that captures this relational nature of the world is that of a graph or, as we also call it, a network. We deal with a multitude of networks every day, from using our metabolic and neural networks to properly function as biological organisms, to navigating through online social networks communicating with family and friends, to our financial institutions' making decisions upon strategic partnership and competing with their international peers. Due to the deep embeddedness of networks into our lives, it is natural to expect that an individual scientific discipline would arise, aiming at understanding and improving our interconnected world. Today, we call this discipline *network science* [1], that has recently emerged as an independent discipline, having amalgamated the parts of numerous well-established fields—such as graph theory, statistical physics, computer science, game theory, control theory, behavioral economics, biology and neuroscience, sociology and social psychology—viewing the world through the abstraction of graphs.

The field of network science has two major sub-fields: the well-established *static network science* that studies networks not changing in time; and a much younger *dynamic network science* that deals with the networks whose structure and “content” evolve. In what follows, we

will briefly review both above mentioned subfields, delineate the location of network process studies in the field, and describe this thesis' contributions to it.

1.1 Static Network Science

Historically, graph theory can be considered the foundation of static network science, with Euler's Seven Bridges of Königsberg [2] being one of the earliest problems that epitomizes the nature of this subfield: we analyze basic network primitives—nodes, edges, paths and cycles, and, more generally, subgraphs—to make a qualitative or quantitative conclusion about the nature of a network or a group of such. The analysis of static networks typically involves node and edge centrality [3, 4, 5, 6, 7, 8, 9, 10] and clustering [11, 12, 13] properties, as well as reachability considerations when dealing with paths. Below, we review more advanced branches of static network science.

Dissecting Networks: A natural sequel to studying network paths is the investigation of how to optimally split network nodes being guided by the network's structure. This is typically referred to as clustering, having the graph cut optimization problem [14, Sec. A 2.2] at its core. While the problem of optimal cut search has a combinatorial nature, it has a beautiful connection to spectral graph theory [15], which, among other things, provides techniques for approximating optimal cuts relying on eigenvalues and eigenvectors of certain matrices associated with the network [16, 17, 18, 19]. Two applications closely related to clustering are community detection [20, 21] and outlier detection [22].

Pattern Mining: Another common problem is, given a static network, to search it for structural patterns satisfying some desirable properties. For example, such patterns may structurally resemble a prototype network indicating a structural anomaly [23], or we may be interested in

the most frequent subgraph search [24, 25].

Extremal Problems on Networks: This subfield includes topics from extremal graph theory, concerned with either establishing existence, counting, or designing networks with desirable properties. Turán theorem [26] and matrix-tree (Kirchhoff’s) theorem are two examples of extremal graph theory results. Additionally, we are often interested in the discovery of such extremal properties or substructures of a given network as shortest paths [27, 28, 29, 30] and optimal network flows [31, 32, 33].

Analysis of Network Sets and Sequences: More generally, we may be provided with a set or a series of networks—being unaware of their origin—and aim at performing analyses in a space comprised of these networks. For example, having abstracted chemical compounds as networks, we may be interested in classifying them into safe and unsafe for use in drug composition. Similarly, we may want to cluster sets of networks; design predictive models where each sample in a supervised learning task is a network; detect anomalies in and extrapolate time series of networks; and search through and index network databases. Two groups of methods of crucial importance for this subfield are graph matching and graph embedding methods, which target assessing similarity between networks and mapping networks into vector spaces, respectively [34, 35, 36, 37, 38].

Extension to Temporal Networks: Very recently, some of the problems that have traditionally been studied in the context of static networks started to be translated to the case of dynamic networks. These include the recent works on temporal pattern mining [39], temporal graph cuts [40], and temporal centrality measures [41, 42, 43].

1.2 Dynamic Network Science

Dynamic network science has evolved over the past 60 years from the early works on random graphs and cellular automata to complex models for brain, socio-economic, financial, and technological networks. The dynamic network science has two natural constituents, focusing on the evolution of either network structure or both the structure and the nodes' states.

Dynamic Network Formation: The fundamental questions in dynamic network formation are how networks are formed, and, given the local laws of network structure formation, what global phenomena the network formation process leads to. One type of results within this category are the random graph models—commonly studied in statistical physics and complex networks communities—with notable examples of Erdős-Rényi [44], Barabási-Albert [45], and Watts-Strogatz [12] models. Another type of results here are the models of strategic network formation, that deal with systems comprised of intelligent agents making decisions, driven by a combination of long- and short-term local and system-wide utility maximization goals. A representative set of examples is provided by Jackson [46, Ch. 6]. Such results are often encountered in economics, game theory, multi-agent systems, and control. While some random graph models also assume agent “interaction”—such as the preferential attachment mechanism underlying Barabási-Albert model—in strategic network formation, the interaction between agents usually reflects some sophisticated socio-economic, physical, or biological laws.

Dynamic Processes in Networks: At the apex of dynamic network science are the studies of dynamic processes occurring in networks [47]. As before, the network's structure can evolve, but, most importantly, the nodes of the network change their states based on the states of other nodes, guided by the network structure. Network processes can be as simple as John Conway's Game of Life [48], and as fundamental as heat diffusing through a complex material. However, this subfield is mostly dealing with higher-level processes of crucial importance to

the well-being of the human society, such as information spread through communication networks, opinion formation in social networks, bio-chemical processes in metabolic and neural networks, virus spread through contact networks, car traffic flow through road networks, failure spread in economic and power networks, as well as coordination in technological networks.

In what follows, we outline the structure of the field of dynamic network processes, and describe how this thesis—having network processes at its core—contributes to the field.

1.3 Network Processes: Analysis, Modeling, Control

Network processes have surrounded people for thousands of years. Heat diffusing through a complex material, warriors coordinating on the battlefield, commodities spreading through a network of merchants, people passing rumors via word of mouth, innovations proliferating through an organizational network—whenever we see a system of interacting entities, we deal with a network process. This thesis is biased towards the processes occurring in social networks, as such processes have a drastic impact upon people’s social well-being. Recently, with the emergence of massive-scale online social networks, as well as collaboration platforms and tools, studying network processes has become feasible and important as never before.

The *central theme* of this thesis is *analysis, modeling, and control of dynamic processes in networks*, with an emphasis on social networks. In this section, we delineate the scope of the inter-disciplinary field of dynamic network processes along three facets:

1. **Analysis**—creation of scalable methods for model-driven analysis of observed network processes.
2. **Modeling**—design and theoretical analysis of plausible complex, yet tractable models for network processes based on the knowledge about the process’ domains.

3. **Control**—development of theories and the accompanying scalable algorithms to guide network processes for optimizing and/or safeguarding their performance.

Below, we discuss each of the above mentioned aspects of network process studies, and summarize this thesis' contribution along each of them.

Analysis: Many network processes can be at least partially observed, and we may be interested in making some qualitative or quantitative conclusions about the process' evolution based on these observations. For example, given a power grid, and having measured electric current flow through a subset of edges at multiple time points, we may want to detect and spatially localize failures in the grid. As another example, having observed how opinions of users of a social network evolve over time, we may want to either detect when the opinions evolve anomalously—such as being a result of a viral marketing campaign—or predict how opinions of the users will evolve in the future. It is important to notice that both above cases require our analysis to be model-driven, that is, our conclusions regarding the process' evolution must be derived with respect to an underlying model that describes the process' expected behavior.

Chapter 2 of this thesis is dedicated to the design of one such model-driven method for the analysis of opinion evolution in social networks. While designed with an eye on applications to social networks, our results generally extend to other domains where processes can be described by probabilistic models of the same type as the models for opinion dynamics.

Modeling: Scientists have been modeling complex processes for centuries. For example, one of the classical physical models—the heat flow equation $\dot{x} = \alpha\Delta x$ having a straightforward discrete-space analog $\dot{x} = -\alpha Lx$, with L being the network's Laplacian—was formulated by Fourier more than 200 years ago. Since then, the need for modeling has diffused from physical systems to social, economic, biological, and technological systems, yet, the complexity of the network models from these new domains is much higher, proportionally to the complexity

of the considered systems, such as the human society or the brain. In this thesis, under the category of network process modeling, we subsume the works targeting design and analysis of non-linear models for network processes that would simultaneously be rich enough to capture the complexity of the underlying domains, and tractable, allowing derivation of analytic insights into the nature of the network processes.

Chapter 3 of this thesis is dedicated to the design of a general non-linear model capturing complex dynamics of polar opinion formation in social networks. This model is a non-linear version of Laplacian flow, and besides capturing opinion formation behavior, is also connected to the process of non-linear heat flow. In addition to constructing this general model and its several specialized versions connected to existing socio-psychological theories, we describe an approach towards the analysis of such non-linear models that can be reused for the design of other complex models from non-social domains.

Control: Having acquired understanding of a network process through its modeling and analysis, the next natural step is to control that process, aiming at either optimization of that process' performance or maintenance of the latter in the face of endogenous or exogenous attacks or shocks incurred by the process. The notion of control is hardly new, and has been actively studied as of 19'th century by applied mathematicians and physicists. Nowadays, control theory is a well-established field, interwoven into the related engineering disciplines typically dealing with physical or technological networks. More recently, mechanism design theory emerged within the fields of economics and game theory, targeting optimization of processes in economic and financial networks. Further, traditional methods of control theory are now being adapted to studying processes in biological networks.

In Chapter 4 of this thesis, we reflect upon the notion of control applicably to social systems, and design one mechanism that can efficiently nullify exogenous attacks upon the process of user opinion formation in social networks via strategic network augmentation, making that

process uncontrollable. Our theoretical and algorithmic insights contribute to the emergence of social control theory within the general area of network process control.

Having outlined the central subject studied in this thesis, in the following section, we provide an overview of the present thesis.

1.4 Overview of the Present Work

This thesis consists of three major parts, contributing to the analysis, modeling, and control dimensions of network process studies, respectively.

Analysis: In Chapter 2, we design a general scalable method to estimate the likelihood of network process’ transitioning between states in the context of opinion dynamics in social networks. This method—a distance measure for network process states termed Social Network Distance (SND)—computes the (log-)likelihood of the most likely transition between two observed user opinion distributions under an opinion dynamics process in a large online social network. SND reduces the state transition likelihood definition to that of a transportation problem or, alternatively, a well-known Earth Mover’s Distance. The latter, roughly, computes the “cost” of optimal way to reshape one (opinion) distribution into another one, where the elementary transform—in the edit distance sense—is the transportation of an opinion between two users along the pathways defined by the social network and the chosen opinion dynamics model. While the direct solution of the obtained transportation problem would have super-cubic time-complexity, SND exploits the special structure of the transportation problem and effectively uses a combination of radix-Fibonacci heap-based Dijkstra algorithm and bi-push minimum-cost network flow algorithm for unbalanced bipartite graphs, resulting in SND’s computability in time linear in the network size. In experiments with Twitter data, SND has shown to be effective at detecting controversial events—such as the introduction of “Oba-

macare” in the US in 2010 —that are known to have polarized the US society. SND has also shown to be effective at predicting future opinions of a subset of network users.

Modeling: One fundamental model of opinion formation—stemming from Festinger’s social comparison and cognitive dissonance theories—is DeGroot model stating that members of a social network—whose adjacency matrix’ entries reflect the relative amounts of inter-personal trust—form their opinions via weighted averaging with the opinions of their network neighbors. One limitation of this model is that every person manifests the same opinion adoption behavior, while in reality people have different receptiveness to persuasion. Within the realm of linear models, this drawback is resolved by Friedkin-Johnsen model that allows the network’s members to have different (though, constant) degrees of “stubbornness”.

In Chapter 3, we propose a general non-linear opinion formation model, having incorporated dynamic opinion-dependent susceptibility to persuasion into the opinion formation process¹. This model is particularly suitable for the case of polar opinions, when people’s shifting towards opinion extremes makes them less or more receptive to opinion change. For example, a person shifting towards the Republican ideology is harder to persuade to invert her or his political stance. Alternatively, in a society with strong social norms, people holding conservative opinions following the norms are harder to persuade, while extreme opinions are volatile. The main theoretical question here is how the user opinions will evolve in a long term. We have answered this question for our general non-linear model and its multiple specialized versions, having brought up the tools from non-smooth analysis, and used them together with min-max Lyapunov functions and the suitable version of LaSalle Invariance Principle. One qualitative finding of this analysis states that, as long as a user is at least to some extent susceptible to persuasion, that specific extent does not affect the long-term opinion adoption behavior of that

¹The origins of the idea that susceptibility to persuasion is a function of the held opinion can be traced back at least to Abelson’s 1964 work “Mathematical models of the distribution of attitudes under controversy” [49], and has been periodically revisited since then by Friedkin and others. However, a theoretically analyzed network model had been absent prior to our work.

user. Another finding is an analytic expression for the asymptotic user opinion distribution—dependent upon the structure of the network as well as the placement of “extremists” in it, yet independent of the initial opinions of susceptible users.

Control: Since the process of opinion formation is inherently a *network process*, it is lucrative for marketing and political technologists to tap into online social networks and shift the public opinion distribution towards business-imposed objectives. One widely studied method of affecting—or controlling—opinion formation is influence maximization, where the goal is to strategically influence the opinions of select—desirably, influential—users, so that they efficiently distribute these “right” opinions through the network. The society would, however, benefit from safeguarding the opinion formation process from such external control.

To that end, in Chapter 4, we propose a novel problem of disabling external influence upon the opinion distribution in a social network, as well as an efficient method to solve it. We assume that the specific target for the attack is the so-called asymptotic consensus value—the sum of user opinions weighted by these users’ eigenvector centralities. This value is the asymptotic limit for user opinions in DeGroot model, or, more generally, a value to which the opinions of all network users are attracted. We assume that the adversary maximizes the asymptotic consensus value by altering the opinions of some users. We, then, state DIVER—an NP-hard problem of disabling such external influence attempts via strategically recommending a limited number of edges to the network’s users. Relying on the theory of Markov chains, we provide perturbation analysis that shows how eigenvector centrality and, hence, DIVER’s objective function change in response to an edge’s addition to the network. The latter leads to a pseudo-linear-time heuristic for DIVER, that relies on efficient estimation of mean first passage times in a Markov chain.

Chapter 2

Model-Driven Analysis of Polar Opinion Dynamics in Social Networks

In this chapter, we design methods for model-driven analysis of the observed process of opinion dynamics in a social network—the process of crucial importance in today’s life. One type of the desirable analyses we aim to enable is predicting users’ political preference. For such analysis, it is particularly important to be able to analyze the dynamics of competing polar opinions—such as pro-Democrat vs. pro-Republican—since opinion polarization is natural for the domain of politics. While observing the evolution of polar opinions in a social network over time, can we tell when the network evolved abnormally? Furthermore, can we predict how the opinions of the users will change in the future? To answer such questions, it is insufficient to study individual user behavior, since opinions can spread beyond users’ ego-networks. Instead, we need to consider the opinion dynamics of all users simultaneously and capture the connection between the individuals’ behavior and the global evolution pattern of the social network.

The key result of this chapter, that enables answering the above stated questions, is the Social Network Distance (SND)—a distance measure that quantifies the likelihood of evolution of one snapshot of a social network into another snapshot under a chosen model of polar

opinion dynamics. SND has a rich semantics of a transportation problem, yet, is computable in time linear in the number of users and, as such, is applicable to the analysis of large-scale online social networks. In our experiments with synthetic and Twitter data, we demonstrate the utility of our distance measure for anomalous event detection. It achieves a true positive rate of 0.83, twice as high as that of alternatives. When used for opinion prediction in Twitter data, SND’s accuracy is 75.6%, which is 7.5% higher than that of the next best method.

The chapter is organized as follows. In Section 2.1 we motivate and informally define the problem. Section 2.2 provides necessary preliminaries, followed by the review of existing distance measures in Section 2.3. The next Section 2.4 is dedicated to the design of the core concept of this chapter—the Social Distance Measure (SND), with the following Section 2.5’s targeting generalization of the Earth Mover’s Distance (EMD)—being at the core of SND—to make it fit the comparison of states of a social network. In Section 2.6, we provide a pseudo-linear method for computation of SND. After experimental results in Section 2.7, we discuss limitations 2.8 of SND. We conclude with a discussion and a summary of the results of this chapter in Section 2.9, and point out potential directions for future research in Section 2.10.

2.1 Introduction

Analysis of opinion formation in the society plays an important role in today’s life. Businesses are interested in advertising their products in social networks relying on viral marketing. Political strategists are interested in predicting an election outcome based on the observed sentiment change of a sample of voters. Mass media and security analysts may be interested in a timely discovery of anomalies based on how a social network “behaves”. Thus, it is important to enable modeling and prediction of user opinion evolution in a social network.

How can we quantify the change in opinions of users with respect to their expected behavior in a social network? How can we predict how the opinions of individual users will evolve in the

future? Having observed the evolution of user opinions over time, can we tell when the opinions evolved abnormally? To answer such questions, we need a distance measure for the comparison of states of a social network that explicitly models user opinion evolution, incorporating both the distribution of user opinions at two time instances and the network structure that defines the pathways for opinion dissemination. In this chapter, we design such a distance measure and employ it for anomaly detection and opinion prediction.

While the dynamics of a social network can be characterized by evolution of both the network’s structure and user opinions, in this chapter we focus on the latter. We assume that there are two *polar opinions* in the network, *positive* and *negative*¹. Users having no or an unknown opinion are termed *neutral*, while those expressing opinion—*active*. A *network state* is comprised of the opinions of all network users at a given time. Polar opinions *compete* in that users are less willing to spread opinions different from their own, yet, are more eager to spread “friendly” opinions. Such competition may arise when the notions the opinions relate to, such as political parties or smartphone brands, are inherently competing.

Having observed the behavior of social network users over time and quantified their opinions, we obtain a time series of network states. Its analysis—whether anomaly detection or extrapolation—is, however, problematic, as network states do not naturally belong to any vector space, and the numerous existing time series analysis techniques cannot be readily applied. Our approach is to treat network states as members of a metric space induced by a distance measure governed by both the network’s structure and user opinions. We propose a semantically and mathematically appealing, as well as efficiently computable distance measure *Social Network Distance (SND)* for the social network states containing polar opinions and demonstrate its utility in two applications. First, we detect which network states in a series are anomalous with respect to the expected opinion evolution, where the latter is determined by a chosen model of polar opinion dynamics. Second, we predict unknown opinions of individual users in

¹Our results can be straightforwardly generalized to the case of any finite discrete polar opinion scale.

a partially observed network state based on the historical dynamics of other users' opinions.

In this chapter, we make the following specific *contributions*:

- We propose SND—the first distance measure suitable for the comparison of social network states containing polar opinions under a chosen model of opinion dynamics.
- We develop a scalable method for exact computation of SND in time linear in the number of users in the network. This is achieved via exploiting the special structure of the transportation problem underlying SND and the use of advanced shortest path and minimum-cost network flow algorithms.
- We demonstrate the utility of SND in two application with both synthetic and Twitter data. Using SND for anomaly detection, we achieve a true positive rate of 0.83, twice as high as that of alternatives. When used for user opinion in Twitter data, SND's prediction accuracy is 75.6%, which is 7.5% higher than that of the next best method.

2.2 Preliminaries

Prior to designing our distance measure, we provide preliminaries upon which the subsequent design will rely. The summary of notation for this chapter is provided in Table 2.1.

2.2.1 Network and Network States

We are given a social network $G(V, E)$, where V ($|V| = n$) is the set of nodes (users) and E is the set of edges (social ties). At each point in time, each user holds a quantified opinion on the chosen topic of interest. In this chapter, we will quantify the opinions on a discrete scale $\{+1, 0, -1\}^2$, with 0 standing for neutrality, and +1 and -1 corresponding to two polar

²There is a great body of research on the methods for opinion quantification based on user-generated content, including [50, 51]. Our focus is, however, on the analysis of *how opinions spread*, rather than their quantification.

$\mathbb{1}$	vector of all ones
$\mathbb{0}$	vector of all zeros
\otimes	Kronecker product
$\text{diag}(v)$	diagonal matrix with vector v as the main diagonal
$G(V, E)$	network with nodes V ($ V = n$) and edges E ($ E = m$)
$X, X(t) \in \{+1, 0, -1\}^n$	network state (at time t) comprised of all users' opinions
I_P^v	set of users holding opinion v in network state P
$\mathbb{P}_{ij}(P, v)$	likelihood of user j acquiring opinion v from user i in network state P
n_Δ	number of users who changed opinion between two network states

Table 2.1: Summary of notation used in the design of Social Network Distance

alternatives, such as the Democrats vs. the Republicans or iOS vs Android. Note, however, that the results we obtain in Section 2.4-2.6 also hold for a more general case of any finite discrete opinion scale. The opinions of all users at time t comprise *network state* $G(t) \in \{+1, 0, -1\}^n$.

2.2.2 User Opinions and Their Dynamics

We assume that the dynamics of user opinions is governed by a chosen in advance opinion dynamics model \mathcal{M} that provides $\mathbb{P}_{ij}(P, v)$ —the likelihood of user j adopting opinion $P_i = v \in \{+1, -1\}$ of user i in network state P . Below, we provide three example definitions of $\mathbb{P}_{ij}(P, v)$ for three different opinion dynamics models.

Independent Cascade Model: For the version of the Independent Cascade model [52] supporting multiple opinion values, $\mathbb{P}_{ij}(P, v)$ is defined as follows

$$\mathbb{P}_{ij}^{\text{IC}}(P, v) = \begin{cases} 0 & \text{if } d_j(\{i\}) > d_j(I_P^{+1} \cup I_P^{-1}), \\ 1 & \text{else if } P_i = v \text{ and } P_j = v, \\ \frac{\max(0, p_{ij} - \varepsilon)}{\sum_{i \in \text{act}(P, j)} p_{ij}} & \text{else if } P_i = v \text{ and } P_j = 0, \\ \varepsilon & \text{otherwise,} \end{cases}$$

where I_p^v is a set of users holding opinion v in network state P , p_{uv} is an edge activation probability [53], $d_j(I)$ is the length of the shortest path from users I to user j , $act(P, j) = \{k \mid k \in I_p^{+1} \cup I_p^{-1} \text{ and } d_j(\{k\}) = d_j(I_p^{+1} \cup I_p^{-1})\}$ is the set of users active in network state P closest to user j , and ε is a negligible likelihood of an “impossible” event.

In the original model, $\varepsilon = 0$, that is, neutral users cannot spread positive or negative opinions to others, and active users neither drop their opinions nor spread opinions opposite to their own. That, however, would lead to the distances between many network states to be $+\infty$ since the opinion evolution may follow the assumed opinion dynamics model *not exactly*. In order to obviate this issue, we aim to—instead of just declaring two network states qualitatively unreachable—to always quantify the distance between them, and, thus, assign some negligible probabilities ε to the events that original opinion dynamics models posit as impossible.

Linear Threshold Model: For the version of Linear Threshold model [54] supporting multiple opinion values, $\mathbb{P}_{ij}(P, v)$ is defined as

$$\mathbb{P}_{ij}^{\text{LT}}(P, v) = \begin{cases} 0 & \text{if } i \notin N^{\text{in}}(P, j), \\ 1 & \text{else if } P_i = v \wedge P_j = j, \\ \frac{(1-\varepsilon)\omega_{ij}}{\sum_{k \in N^{\text{in}}(P, j)} \omega_{kj}} & \text{else if } P_i = v \text{ and } P_j = 0 \text{ and } \sum_{k \in N^{\text{in}}(P, j)} \omega_{kj} \geq \theta_j, \\ \varepsilon & \text{otherwise,} \end{cases}$$

where ω_{ij} is an edge weight reflecting relative influence of i upon j , θ_i is user i 's opinion switching threshold, $N^{\text{in}}(P, j)$ is the set of j 's in-neighbors active in network state P , and ε has the same semantics of a negligible likelihood of an “impossible” event as in the earlier case of the Independent Cascade model.

Basic Model: Besides the given above definitions of $\mathbb{P}_{ij}(P, v)$ for two well-established models, we also provide a definition for a simpler intuitive model that will be used in our experiments. According to this model, $\mathbb{P}_{ij}(P, v)$, translated into log-likelihoods, are defined as

$$-\log \mathbb{P}_{ij}(P, v) = \begin{cases} c_{adverse} & \text{if } P_i = -v \text{ or } P_j = -v, \\ c_{neutral} & \text{else if } P_i = 0, \\ c_{friendly} & \text{else if } P_i = v \text{ and } P_j \neq -v, \end{cases} \quad (2.1)$$

where $c_{adverse}, c_{neutral}, c_{friendly} \in \mathbb{R}^+$ are constant log-likelihoods of adopting adverse, neutral, or friendly opinion (relatively to opinion v), respectively. Thus, the basic model assumes that users willingly spread opinions similar to their own ($c_{friendly}$ is small), are unwilling to spread adverse opinions ($c_{adverse}$ is large), with the behavior of neutral users being somewhere in-between ($c_{friendly} < c_{neutral} < c_{adverse}$).

2.2.3 Earth Mover's Distance

At the core of our approach towards analyzing observed opinion dynamics is the computation of distances between network states $P, Q \in \{+1, 0, -1\}^n$ using a distance measure, whose semantics will be similar to that of one well-studied distance measure—Earth Mover's Distance (EMD) [55]. Originally defined as an edit-distance for histograms (which, in our case, can be seen simply as n -dimensional non-negative vectors), EMD computes the cost of an optimal transformation of one histogram into another, where an elementary edit operation is transportation of a unit of mass from one histogram bin to another bin. The costs of these elementary transforms are collectively referred to as the *ground distance*.

Formally, EMD between two histograms $P \in \mathbb{R}^+{}^n$ and $Q \in \mathbb{R}^+{}^m$ (that, in our case, will be derived from network states) over ground distance $D \in \mathbb{R}^+{}^{n \times m}$ (that, in our case, will be

defined based on the distances between users in the network and the likelihoods of the opinions spreading between them) is the solution to the problem of optimal mass transportation from suppliers $\{P_i\}$ to consumers $\{Q_j\}$ with respect to transportation costs $\{D_{ij}\}$:

$$\begin{aligned} \text{EMD}(P, Q, D) &= \sum_{i=1}^n \sum_{j=1}^m D_{ij} \hat{f}_{ij} / \sum_{i=1}^n \sum_{j=1}^m \hat{f}_{ij}, & (2.2) \\ \{\hat{f}_{ij}\} &= \arg \min_{\{f_{ij}\}} \sum_{i=1}^n \sum_{j=1}^m f_{ij} D_{ij}, \quad \sum_{i=1}^n \sum_{j=1}^m f_{ij} = \min \left\{ \sum_{i=1}^n P_i, \sum_{j=1}^m Q_j \right\}, \\ f_{ij} &\geq 0, \quad \sum_{j=1}^m f_{ij} \leq P_i, \quad \sum_{i=1}^n f_{ij} \leq Q_j, \quad (1 \leq i \leq n, 1 \leq j \leq m), \end{aligned}$$

where $\{\hat{f}_{ij}\}_{n \times m}$ is an optimal solution or transportation plan.

In addition to having the semantics that will suit our distance measure design goals, EMD is metric, as the following theorem states.

Theorem 1 (Metricity of EMD [55]). *If all network states (network states) under comparison have equal total masses, and the underlying ground distance is metric, then EMD is metric.*

Metricity of EMD will allow our own EMD-based distance measure to also be metric, which—in addition to making a distance measure “natural”—can be exploited to improve practical performance of distance measure-based algorithms in applications [56].

2.3 Related Work

Given our distance measure-driven approach towards analysis of the network process of opinion dynamics, the body of works directly related to ours describes various distance measures. In this section, we review some representative distance measures and point out why none of them is suitable for the model-aware analysis of the process of polar opinion dynamics, which necessitates the design of our Social Network Distance in Section 2.4.

There is a large number of existing distance measures used in finite-dimensional vector spaces, including ℓ_p , Hamming, Canberra, Cosine, Kullback-Leibler, and Quadratic Form [57] distances. However, none of them are suitable for the comparison of social network states, since these distance measures either compare opinion vectors coordinate-wise, thereby, not capturing user interaction in the network, or in the case of Quadratic Form distance, capture the user interaction in a very limited way, being unaware of the underlying process that causes the difference between two user opinion distributions.

Existing graph-oriented distance measures are also unsuitable for comparing network states with polar opinions. The first type of such distance measures is graph isomorphism-based distance measures, such as *largest common subgraph* [58]. They are node state-oblivious, and, hence are not applicable to the comparison of network states. Another type of graph distance measures is *Graph Edit Distance (GED)-based* distance measures [59] that define the distance between two networks as the cost of an optimal sequence of node or edge insertions, deletions, and substitutions, transforming one network into another. GED can be node state-aware, but its value is not interpretable from the opinion dynamics point of view, and even its approximate computation takes cubic time (a single computation of GED on a 10k-node network on our hardware takes about a month). DELTACON [60] is a scalable graph-oriented distance measure, yet, it quantifies the networks' structural difference, while we focus on node states.

A third class of distance measures includes *iterative distance measures* [61, 62, 63, 64], which express similarity of the nodes of two networks recursively, use a fix-point iteration to compute node similarities, and, then, aggregate node similarities to obtain the similarity of two networks. These share the problem of GED—they do not capture the way competing opinions spread in the network. The same drawback is shared by the related diffusion-based distance measures [65, 66, 67], that compare network states by quantifying how differently a heat diffusion process proceeds in the network when the network states defined the initial temperatures of the nodes.

The last class includes *feature-based* distance measures [68, 69, 70], which compare either the distributions of local node properties (e.g., degree, clustering coefficient) or the spectra of two networks. Despite their efficient computability, such distance measures do not fit the comparison of network states with polar opinions. The spectral distance measures are inadequate because they do not deal with node states directly³, while other feature-based distance measures only deal with summaries based on opinion of each kind, thus, being unable to capture competition of polar opinions.

2.4 Social Network Distance

The *central problem* we address in this chapter is as follows:

Problem. *Given two network states $P, Q \in \{+1, 0, -1\}^n$ and assuming that the dynamics of user opinions is governed by model \mathcal{M} , define and compute the distance between network states P and Q reflecting how likely either of these network states has evolved into another under model \mathcal{M} .*

According to the problem definition above, if network state Q has evolved from network state P closely following model \mathcal{M} , then the resulting distance should be small; if, however, either network states P and Q are unrelated, or they are related, but the opinions evolved following a model very different from \mathcal{M} , then the distance should be large. Next, we will, first, translate the above defined problem into the formal language of probability, and, after that, address the question of making the obtained formalization tractable.

Let us put $X(t)$ to be the state of the network at time t , and $X(1, \dots, k) = X(1), \dots, X(k)$ to be a *network state evolution path*—a sequence of states over which the network evolved from

³Even if node states are artificially encoded into a network’s structure, there is still a possibility for two structurally different networks to be isospectral, that is, to have identical spectra and, hence, a zero spectral distance.

time $t = 1$ to time $t = k$. Then, the likelihood of $X(1, \dots, k)$ is defined as

$$\mathbb{P}\{X(1, \dots, k)\} = \prod_{t=2}^k \mathbb{P}\{X(t) \mid X(t-1)\}.$$

Using this notation, a *perfect distance measure* can be defined as

$$d_{\mathcal{M}}(P, Q) = \sum_{k \geq 2} \sum_{\substack{X(1, \dots, k) \\ X(1)=P, X(k)=Q}} \mathbb{P}\{X(1, \dots, k)\}. \quad (2.3)$$

According to its definition, $d_{\mathcal{M}}(P, Q)$ measures the likelihood of user opinions evolving from state $X(1) = P$ to state $X(k) = Q$ along all possible network state evolution paths—as illustrated in Figure 2.1—where the opinion change likelihoods are determined by model \mathcal{M} .

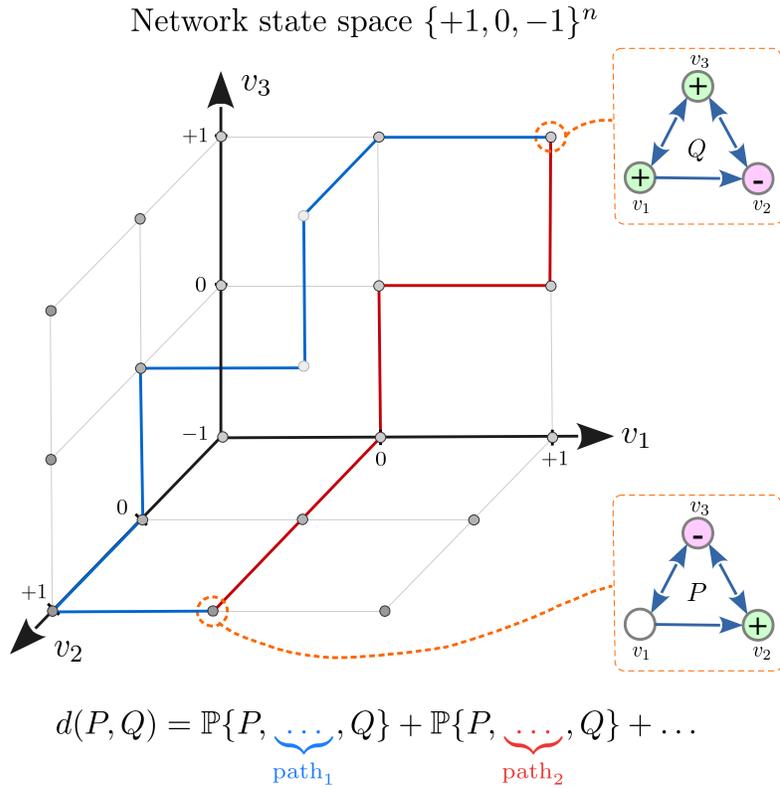


Figure 2.1: $d_{\mathcal{M}}(P, Q)$ accumulates likelihoods of all possible network state evolution paths.

Despite the attractive semantics of $d_{\mathcal{M}}$, its computation is clearly unfeasible, as the number of possible evolution paths between two network states is exponential in the number of nodes. To come up with a tractable alternative, we simplify $d_{\mathcal{M}}$ by making several assumptions.

Assumption 1 (Maximum-likelihood opinion evolution). *Opinions evolve along to the most likely network state evolution path.*

According to Assumption 1—standard for maximum likelihood estimation methods—we will not care about every possible network state evolution path, instead, focusing on the most likely one. Under this Assumption $d_{\mathcal{M}}$ simplifies to $d^{(1)}$ as follows:

$$d^{(1)}(P, Q) = \max_{k \geq 2} \max_{\substack{X(1, \dots, k) \\ X(1)=P, X(k)=Q}} \mathbb{P}\{X(1, \dots, k)\}. \quad (2.4)$$

To further simplify the obtained distance measure, we will target the term being maximized in (2.4), and make another assumption.

Assumption 2 (Independent Markovian opinion acquisition). *Users acquire opinions asynchronously, independently of other users' opinion acquisition, relying only on the opinions of network users at the previous time.*

According to this assumption, opinions of different users evolve at individual time scales (in contrast to synchronous opinion evolution, where all users simultaneously update their opinions), and depend only on the previously observed opinions of other users. This assumption simplifies distance measure $d^{(1)}$ into $d^{(2)}$ as follows.

$$d^{(2)}(P, Q) = \max_{k \geq 2} \max_{\substack{X(1, \dots, k) \\ X(1)=P, X(k)=Q}} \prod_{i=1}^n \prod_{t=2}^k \mathbb{P}\{X_i(t) \mid X(t-1)\} \quad (2.5)$$

Now, to make the maximization task in (2.5) tractable, we will make one more assumption.

Assumption 3 (Opinion source uniqueness). *An opinion is adopted by a user from a single most likely source.*

This assumption—previously used by Gomez-Rodriguez et al. [71] in the context of cascade inference—is natural in those cases when obtaining knowledge immediately causes or is equivalent to opinion acquisition, such as in the case of learning an incriminating piece of news about a politician. In these situations, a contact with a single information source is sufficient to acquire opinion, and contacting additional sources would not solidify that opinion even further.

If we put $f_{ji} \in [0, 1]$ to be the likelihood of user j being the source for opinion acquisition by user i , $d^{(2)}$ simplifies under Assumption 3 into $d^{(3)}$ as follows:

$$d^{(3)}(P, Q) = \prod_{v \in \{+1, -1\}} \max_{\{f_{ji}\}} \prod_{i \in I_Q^v} \sum_{j \in I_P^v} f_{ji} \mathbb{P} \left\{ \begin{array}{l} j \text{ infects } i \text{ with opinion } v \text{ along the most likely} \\ \text{evolution path } P = X(1), X(2), \dots, X(\cdot) = Q \end{array} \right\}, \quad (2.6)$$

$$\sum_j f_{ji} = |Q_i|, \quad \sum_i f_{ji} \in \mathbb{Z}^+,$$

where I_P^v is the set of users holding opinion v in network state P . Here, f_{ji} can alternatively be viewed collectively as a probabilistic mapping between opinion sources and destinations. The obtained distance measure (2.6) is equivalent to the following expressed using log-likelihoods

$$d^{(4)}(P, Q) = \sum_{v \in \{+1, -1\}} \min_{\{f_{ji}\}} \sum_{i \in I_Q^v} \sum_{j \in I_P^v} f_{ji} \left(-\log \mathbb{P} \left\{ \begin{array}{l} j \text{ infects } i \text{ with opinion } v \text{ along the most likely} \\ \text{evolution scenario } P = X(1), X(2), \dots, X(\cdot) = Q \end{array} \right\} \right)$$

$$= \sum_{v \in \{+1, -1\}} \min_{\{f_{ji}\}} \sum_{i \in I_Q^v} \sum_{j \in I_P^v} f_{ji} D_{ji}(P, v), \quad (2.7)$$

$$\sum_j f_{ji} = |Q_i|, \quad \sum_i f_{ji} \in \mathbb{Z}^+. \quad (2.8)$$

In the expression above, $D_{ji}(P, v) \in \mathbb{R}^{+n \times n}$ is the log-likelihood (or cost) of opinion v 's spreading from user j to user i along the most likely path through the network in state P . Provided that for each pair of nodes j and i , the chosen opinion dynamics model \mathcal{M} defines the likelihood

$\mathbb{P}_{ji}(P, v)$ of opinion v spreading through edge (j, i) in network state P —as per (2.1)— $D_{ji}(P, v)$ is defined as *the length of the shortest path from node j to node i in the network whose structure is identical to that of the network that P is defined over, and whose edge (j, i) is weighted with $-\log \mathbb{P}_{ij}(P, v)$.*

Finally, if we relax $f_{ji} \in [0, 1]$ in (2.7) to be arbitrary non-negative reals, the obtained distance measure (2.7)-(2.8) almost exactly matches Earth Mover’s Distance (EMD) described in detail in Section 2.2.3:

$$\begin{aligned} \text{EMD}(P, Q, D) &= \sum_{i=1}^n \sum_{j=1}^m D_{ij} \hat{f}_{ij} / \sum_{i=1}^n \sum_{j=1}^m \hat{f}_{ij}, \\ \{\hat{f}_{ij}\} &= \arg \min_{\{f_{ij}\}} \sum_{i=1}^n \sum_{j=1}^m f_{ij} D_{ij}, \quad \sum_{i=1}^n \sum_{j=1}^m f_{ij} = \min \left\{ \sum_{i=1}^n P_i, \sum_{j=1}^m Q_j \right\}, \\ f_{ij} &\geq 0, \quad \sum_{j=1}^m f_{ij} \leq P_i, \quad \sum_{i=1}^n f_{ij} \leq Q_j, \quad (1 \leq i \leq n, 1 \leq j \leq m). \end{aligned}$$

Besides the difference in the the normalization factor $\sum_{i,j} \hat{f}_{ij}$, EMD imposes an extra constraint upon the sum of f_{ij} , requiring $\sum f_{ij} = \min\{\sum_i P_i, \sum_j Q_j\}$. For us, it would roughly mean that the number of active users—holding opinions $+1$ and -1 —in both network states P and Q should be equal. The latter, however, does not hold in practice—as shown in Figure 2.2—as already active users may be unwilling to become neutral, and even more active users may appear while information spreads through the network.

This difference will disappear when we replace EMD with its generalization EMD* in Section 2.5.

Based on the obtained expression (2.7) and expression (2.2) for EMD, we can define the nonsymmetric version of our Social Network Distance as follows

$$\text{SND}^{\text{nonsym}}(P, Q) = \sum_{v \in \{+1, -1\}} \text{EMD}(P^v, Q^v, D(P, v)),$$

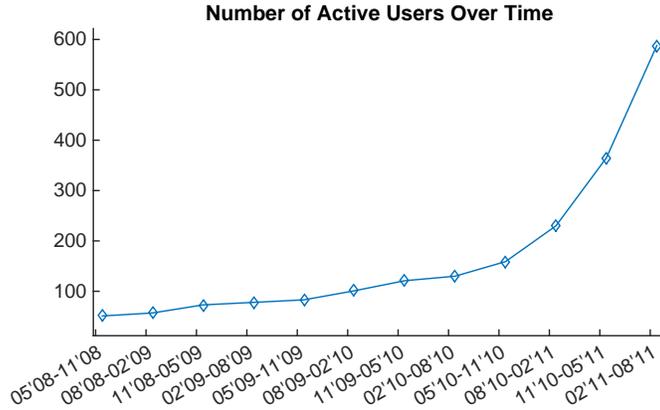


Figure 2.2: Growth of the number of active users—holding non-neutral opinions with respect to topic “Obama”—in a sample of Twitter network over time.

where P^v (Q^v) is a network state P (Q) in which all the users holding opinions different from v are considered neutral, and ground distance $D(P, v)$ consists of log-likelihoods as defined in (2.7). The defined above $SND^{nonsym}(P, Q)$ is suitable for the comparison of time-ordered network states—when we know that network state P preceded network state Q in time. However, we would like to be able to compare arbitrary unordered network states. This may be important for such applications as nearest neighbor network state search, where nearest neighbor candidate network states may either precede or succeed in time the target network state. To that end, instead of building upon $d_{\mathcal{H}}(P, Q)$, we will be interested in $\sqrt{d_{\mathcal{H}}(P, Q) \cdot d_{\mathcal{H}}(Q, P)}$, and, having repeated all the simplifications with the obtained expression, define our Social Distance Measure as follows.

Definition 1 (Social Network Distance (SND)).

$$SND(P, Q) = \frac{1}{2} \sum_{v \in \{+1, -1\}} \left[EMD(P^v, Q^v, D(P, v)) + EMD(Q^v, P^v, D(Q, v)) \right], \quad (2.9)$$

where EMD is a version of Earth Mover’s Distance. Notice that SND is a linear combination of multiple instances of EMD , so the following theorem trivially holds.

Since SND in the definition above is a linear combination of multiple instances of EMD, it is metric provided that the underlying EMD is such.

Theorem 2 (Metricity of SND). *SND is metric as long as the underlying EMD is metric.*

To summarize, we have defined Social Network Distance (SND)—a distance measure that approximates the (log-)likelihood of a network state’s most-likely transition into another network state. As we have mentioned earlier, SND was defined via EMD ignoring the normalization factor in the definition (2.2) of EMD as well as the difference in the constraints between (2.2) and (2.8). In the following Section 2.5, we will generalize EMD to address both these issues, and use its generalization $\widehat{\text{EMD}}$ in definition (2.9) of SND.

2.5 Generalized Earth Mover’s Distance

2.5.1 Why Do We Need a New Earth Mover’s Distance?

Our distance measure SND defined in the previous Section 2.4 uses an EMD, such as the original Earth Mover’s Distance [55], as a building block. Unfortunately, the original EMD cannot adequately compare network states P and Q having different total masses, that is, network states with $\sum P_i \neq \sum Q_j$ —it ignores the mass mismatch $|\sum P_i - \sum Q_j|$. Applicably to states of a social network, that limitation is particularly pronounced, as, usually, subsequent network states have more active users and, hence, a larger total mass than preceding network states.

There are several versions of EMD that address the original EMD’s neglect for network state mass mismatch. One of them, namely, $\widehat{\text{EMD}}$ [72] augments EMD with an additive mass

mismatch penalty as

$$\widehat{\text{EMD}}(P, Q, D) = \text{EMD}(P, Q, D) \min \left\{ \sum_{i=1}^n P_i, \sum_{j=1}^n Q_j \right\} + \alpha \max_{i,j} \{D_{ij}\} \left| \sum_{i=1}^n P_i - \sum_{j=1}^n Q_j \right|,$$

where α is a constant parameter. Term $\alpha \max_{i,j} \{D_{ij}\} |\sum P_i - \sum Q_j|$ represents the mass mismatch penalty that depends only on the magnitude of the mass mismatch and the maximum ground distance, thereby, being unable to capture the fine details of the network's structure that D can depend upon. This is, however, inadequate for the comparison of the states of a social network, because the network's behavior depends not only on the number of new activations, but as importantly on where in the network these new activations occur.

Another Earth Mover's Distance version EMD^α [73] extends each network state with an extra node—the *bank node*—whose value is chosen to equalize the network states' total masses. An example of such an extension is shown in Figure 2.3.

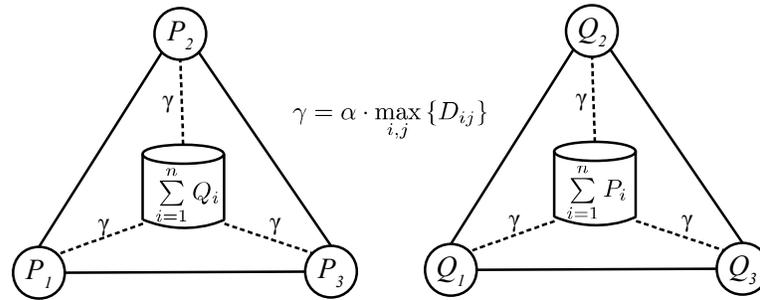


Figure 2.3: Network states P and Q defined over the same network are extended with bank nodes, whose values are chosen, so that the total masses of the extended network states \tilde{P} and \tilde{Q} are equal. The ground distances $\tilde{D}_{bank,i} = \tilde{D}_{i,bank} = \gamma$ from and to the bank node are uniformly defined based on the largest ground distance between the initially present nodes.

Formally, EMD^α is defined as follows.

$$\begin{aligned}
P &= [P_1, \dots, P_n], \quad Q = [Q_1, \dots, Q_n], \\
P_{bank} &= \sum_{j=1}^n Q_j, \quad \tilde{P} = [P_1, \dots, P_n, P_{bank}], \\
Q_{bank} &= \sum_{i=1}^n P_i, \quad \tilde{Q} = [Q_1, \dots, Q_n, Q_{bank}], \\
\tilde{D} &= \left[\begin{array}{c|c} D_{n \times n} & \begin{array}{c} | \\ \alpha \max_{i,j} \{D_{ij}\} \\ | \end{array} \\ \hline \begin{array}{c} - \alpha \max_{i,j} \{D_{ij}\} - \\ \end{array} & 0 \end{array} \right], \\
\text{EMD}^\alpha(P, Q) &= \text{EMD}(\tilde{P}, \tilde{Q}, \tilde{D}) \cdot \left(\sum_{i=1}^n P_i + \sum_{j=1}^n Q_j \right).
\end{aligned}$$

Despite the syntactic differences between EMD^α and $\widehat{\text{EMD}}$, the following Theorem 3 establishes that they are actually numerically equivalent.

Theorem 3. *If ground distance $D \in \mathbb{R}^{n \times n}$ is metric, and $\alpha \geq \frac{1}{2}$, so that both EMD^α and $\widehat{\text{EMD}}$ are metric [73, 72], then $\forall P, Q \in \mathbb{R}^{+n} : \text{EMD}^\alpha(P, Q, D) = \widehat{\text{EMD}}(P, Q, D)$.*

Proof. Without loss of generality, let us assume that $\sum P_i \leq \sum Q_j$, and use the following notation

$$\Delta = \Delta(P, Q) = \left| \sum_{i=1}^n P_i - \sum_{j=1}^n Q_j \right|, \quad \gamma = \alpha \max_{i,j} \{D_{ij}\},$$

so that the expression for $\widehat{\text{EMD}}$ gets rewritten as

$$\widehat{\text{EMD}}(P, Q) = \text{EMD}(P, Q) \min \left\{ \sum_{i=1}^n P_i, \sum_{j=1}^n Q_j \right\} + \gamma \Delta.$$

Our aim is to show that EMD^α has exactly the same expression as $\widehat{\text{EMD}}$ as long as they

both are metric. To do so, let us consider how a unit of mass can be transported from network state P to network state Q , both extended with a bank node, as per the definition of EMD^α .

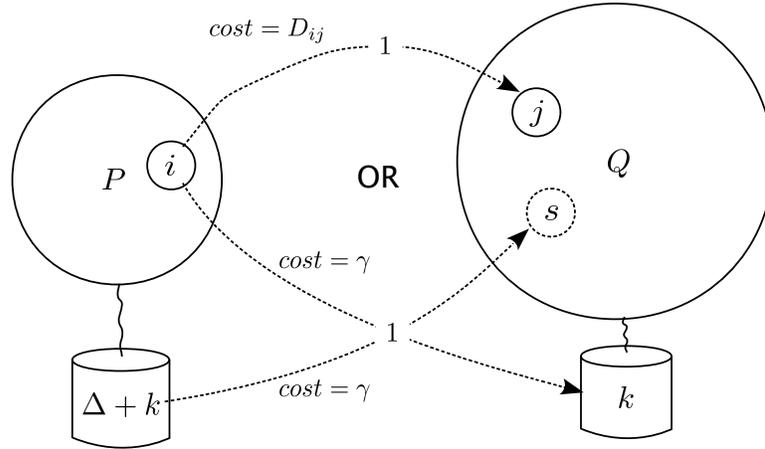


Figure 2.4: Two qualitatively different ways to transport a unit of mass from extended network state $\tilde{P} = [P_1, \dots, P_n, k + \Delta]$ to extended network state $\tilde{Q} = [Q_1, \dots, Q_n, k]$, where $\sum P_i \leq \sum Q_j$. Dashed arrows represent the flow of mass. The bank node is attached to every node of each network state. $k = \sum P_i$, so that the total masses of two extended network states are equal.

As shown in Figure 2.4, there are two qualitatively different alternatives for moving a unit of mass from regular (non-bank) node i of network state \tilde{P} : a unit of mass can be moved either to a regular node j or to the bank node of \tilde{Q} .

In the first case, the total transportation cost of a unit of mass is exactly the ground distance $\tilde{D}_{ij} = D_{ij}$ between regular nodes i and j .

In the second case, the immediate cost of transportation to the bank node is $\tilde{D}_{i,bank} = \gamma$. However, because we have routed mass from a regular node to the bank node, there exists a regular node s in \tilde{Q} having a “mass deficit” that has to be fulfilled from the bank node of \tilde{P} . Thus, if we move a unit of mass from a regular node of \tilde{P} to the bank node of \tilde{Q} , there is an additional incurred cost γ of moving an additional unit of mass from the bank node of \tilde{P} to some regular node of \tilde{Q} . Hence, the total cost of transportation of a unit of mass in the second

case is

$$\gamma + \gamma = 2\alpha \max_{i,j} D_{ij} \geq (\text{since } \alpha \geq 0.5) \geq \max_{i,j} D_{ij}.$$

Thus, from the point of view of optimal mass transportation, it may never be preferable to move a unit of mass from a regular node to the bank node if there is an option to transport mass from a regular node to another regular node. Consequently, an optimal solution to the EMD^α 's transportation problem can be decomposed as follows.

$$\begin{aligned} \text{EMD}^\alpha(P, Q, D) &= \text{EMD}(\tilde{P}, \tilde{Q}, \tilde{D}) \left(\sum_{i=1}^n P_i + \sum_{j=1}^n Q_j \right) = \min_{\{f_{ij}\}} \sum_{i,j=1}^{n+1} f_{ij} \tilde{D}_{ij} = (\text{let } b = n+1) \\ &= \min_{\{f_{ij}\}} \left[\underbrace{\sum_{i,j=1}^n f_{ij} \tilde{D}_{ij}}_{\substack{\text{regular nodes} \\ \text{to regular nodes}}} + \underbrace{\sum_{i=1}^n f_{ib} \tilde{D}_{ib}}_{\substack{\text{regular nodes} \\ \text{to bank node}}} + \underbrace{\sum_{j=1}^n f_{bj} \tilde{D}_{bj}}_{\substack{\text{bank node to} \\ \text{regular nodes}}} + \underbrace{f_{bb} \tilde{D}_{bb}}_{\substack{\text{bank node} \\ \text{to bank node}}} \right] \\ &= \min_{\{f_{ij}\}} \left[\sum_{i,j=1}^n f_{ij} D_{ij} + \gamma \underbrace{\sum_{j=1}^n f_{bj}}_{\Delta} \right] = \min_{\{f_{ij}\}} \left[\sum_{i,j=1}^n f_{ij} D_{ij} + \gamma \Delta \right] \\ &= \min_{\{f_{ij}\}} \left[\sum_{i,j=1}^n f_{ij} D_{ij} \right] + \gamma \Delta = \text{EMD}(P, Q, D) \times \min \left\{ \sum P_i, \sum Q_j \right\} + \gamma \Delta \\ &= \widehat{\text{EMD}}(P, Q, D). \end{aligned}$$

An additional useful observation, formalized below as Corollary 1—that will subsequently show instrumental in the proof of Theorem 4—is that a particular value of k does not matter for EMD^α , since in every optimal solution of its underlying transportation problem, any amount of mass exceeding Δ in the bank node of the lighter network state is transported at cost $\tilde{D}_{\text{bank},\text{bank}} = 0$ to the bank node of the heavier network state. ■

Corollary 1. For network states $P = [P_1, \dots, P_n]$ and $Q = [Q_1, \dots, Q_n]$, and ground distance D ,

if $\sum P_i = \sum Q_j$ and D is metric, then for all $k \geq 0 \in \mathbb{R}^+$, the following holds.

$$EMD \left([P, k], [Q, k], \left[\begin{array}{c|c} D & \omega \\ \hline -\omega & 0 \end{array} \right] \right) = EMD(P, Q, D),$$

where $[X, k]$ is network state X extended with a single bank node with value k and a uniformly defined ground distance $\omega \geq \frac{1}{2} \max_{i,j} D_{i,j}$ to/from the regular nodes of X .

Informally, Corollary 1 states that, for any two network states of equal total mass, we can increase their total masses by an arbitrary value without affecting the EMD between them.

2.5.2 Generalized Earth Mover's Distance

In this section, in response to the inadequacy of the standard EMD for the purpose of social network state comparison, discussed in the previous section, we propose EMD^* —a new version of Earth Mover's Distance, building upon EMD^α 's idea of augmenting network states to even their masses. However, unlike its predecessor, EMD^* extends network states with *multiple local bank nodes*—as shown in Figure 2.5—and distributes the total mass mismatch over all of them, thereby, relating the mass mismatch penalty to the structure of the network, while achieving the total mass equality of the two network states under comparison.

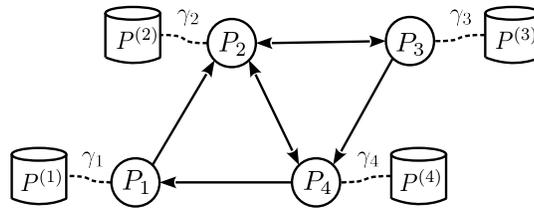


Figure 2.5: A network state $[P_1, \dots, P_4, P^{(1)}, \dots, P^{(4)}]$ extended with *local bank nodes* $P^{(i)}$. The undirected edges to and from the bank nodes—displayed dashed—are weighted with ground distances $\gamma_1, \dots, \gamma_4$, respectively.

Prior to formalizing EMD^* , let us, first, better understand its advantage over the existing EMDs as well as its fitness to the analysis of opinion evolution.

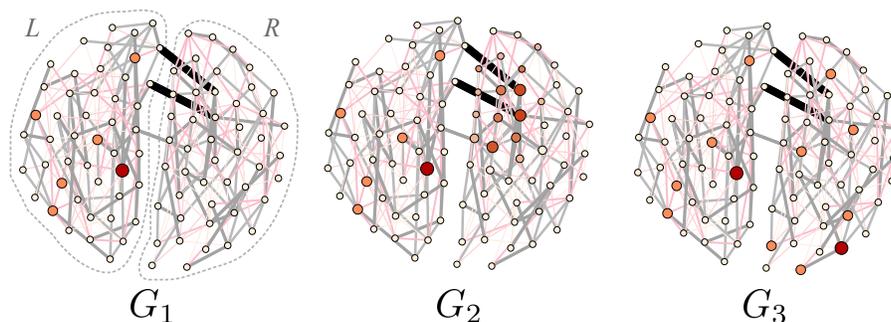


Figure 2.6: Three network states defined over the same two-cluster network.

Consider the example in Figure 2.6. There are three network states defined over the same network, which has two pronounced clusters L and R connected by three bridge edges. The distribution of mass over cluster L is identical in all three states G_i , while cluster R is empty in G_1 and has some differently distributed mass in G_2 and G_3 . In G_2 the extra mass has been “propagated” from cluster L to cluster R through the bridges, while in G_3 the same amount of extra mass has been randomly distributed over cluster R . Thus, if we assume that G_2 and G_3 have “evolved” from G_1 through a (not completely random) network process of mass diffusion, then G_2 should intuitively be closer to G_1 than G_3 is. However, only EMD^* captures this intuition as $\text{EMD}^*(G_1, G_2) < \text{EMD}^*(G_1, G_3)$, while for EMD^α and $\widehat{\text{EMD}}$, G_2 and G_3 are equidistant from G_1 , and for the original EMD, both G_2 and G_3 are identical to G_1 .

We will now turn from the intuition for the formal definition of different components of EMD^* and, eventually, EMD^* itself. The extension of network state $P \in \mathbb{R}^n$ with local bank nodes requires definition of the ground distances γ_i to/from the bank nodes as well as the banks’ values $P^{(i)}$.

Bank node ground distances γ_i : In the extreme case, when $\gamma_i = 0$, in the transportation problem underlying EMD^* , the mass is transported to/from the banks at a zero cost, which

would result in EMD^* 's ignorance of the network state mass mismatch Δ , making it similar to the original EMD up to normalization by Δ . If, on the other hand, γ_i is much larger than the ground distances between regular (non-bank) nodes, then the value of EMD^* will be dominated by the effect of the network state mass mismatch, hiding the impact of the actually present mass. Thus, γ_i should be chosen of the same order as the ground distances D_{ij} between regular nodes, with the particular values of γ_i being empirically learned.

Bank node values $P^{(i)}$: The values of the added bank nodes should be determined based upon two ideas. Firstly, the value of a bank node should intuitively be proportional to the total mass of the node the bank is attached to, thereby, preserving the relative distribution of mass over the network. Secondly, the values of all the bank nodes should be such, that the two network states under comparison have equal total masses. The following definition of value $P^{(i)}$ of a bank node connected to the i 'th node of network state P in the context of comparing network states $P = [P_1, \dots, P_n]$ and $Q = [Q_1, \dots, Q_n]$ incorporates both above mentioned requirements.

$$P^{(i)} = \begin{cases} P_i / (\sum_{j=1}^n Q_j - \sum_{j=1}^n P_j), & \text{if } \sum Q_j > \sum P_j, \\ 0, & \text{otherwise.} \end{cases}$$

Next, we formally define EMD^* . Suppose we are given two network states $P = [P_1, \dots, P_n]$ and $Q = [Q_1, \dots, Q_n]$ defined over a network $G = \langle V, E \rangle$ with ground distance $D_{n \times n}$. Network states P and Q are extended with bank nodes, with $P^{(i)}$ and $Q^{(i)}$ being the values of the bank nodes attached to the i 'th regular node of P and Q , respectively. Ground distances to/from the bank nodes are defined, collectively, as $\gamma = [\gamma_1, \dots, \gamma_n]^\top$.

Then, EMD^* is defined as follows.

Definition 2 (Generalized Earth Mover's Distance (EMD^{*})).

$$\begin{aligned}
 EMD^*(P, Q) &= EMD(\tilde{P}, \tilde{Q}, \tilde{D}) \max \{ \sum P_i, \sum Q_j \}, \\
 \tilde{P} &= [P, P^{(1)}, \dots, P^{(n)}], \quad \tilde{Q} = [Q, Q^{(1)}, \dots, Q^{(n)}], \\
 \tilde{D} &= \left[\begin{array}{c|c} D & D + \mathbf{1}_n \otimes \gamma^\top \\ \hline D + \mathbf{1}_n^\top \otimes \gamma & D + \mathbf{1}_n \otimes \gamma^\top + \mathbf{1}_n^\top \otimes \gamma - 2 \text{diag}(\gamma) \end{array} \right],
 \end{aligned} \tag{2.10}$$

where $P^{(i)}$ is the value of the i 'th bank node that P is extended with (same for Q), $\mathbf{1}_n \in \mathbb{R}^{n \times 1}$ is a vector of all ones, $\text{diag}(v)$ is a diagonal matrix with the elements of vector v on its main diagonal, and \otimes is Kronecker product.

Metricity of EMD^{*}, which can be exploited to improve practical performance of distance-based search in applications [56], is established in the following Theorem 4.

Theorem 4. *Given a finite set \mathcal{H} of network states and metric ground distance D , EMD^{*} defined over D is metric on $\mathcal{H} \times \mathcal{H}$.*

Proof. Let us define $M = \max_{X \in \mathcal{H}} \sum_k X_k < \infty$. Next, we define an auxiliary distance measure EMD' as follows.

$$\begin{aligned}
 EMD'(P, Q, D) &= EMD(P', Q', D'), \\
 P' &= [\tilde{P}, M - \sum \tilde{P}_i], \quad Q' = [\tilde{Q}, M - \sum \tilde{Q}_j], \\
 D' &= \left[\begin{array}{c|c} \tilde{D} & \begin{array}{c} | \\ \max_{i,j} \{ \tilde{D}_{ij} \} / 2 \\ | \end{array} \\ \hline - \max_{i,j} \{ \tilde{D}_{ij} \} / 2 - & 0 \end{array} \right],
 \end{aligned}$$

where \tilde{P} , \tilde{Q} , and \tilde{D} are the extended network states, and the extended ground distance, respectively, as defined by EMD^* . From the definition of EMD^* (2.10), it follows that $\sum \tilde{P}_i = \sum \tilde{Q}_i$ and, hence $M - \sum \tilde{P}_i = M - \sum \tilde{Q}_i = k$. Thus, since $\sum P'_i = \sum Q'_i = M$, D is metric, and $k \geq 0$, from Corollary 1 with $P = \tilde{P}$, $Q = \tilde{Q}$, $D = \tilde{D}$, and $\omega = \frac{1}{2} \max_{i,j} \tilde{D}_{ij}$, we have

$$\begin{aligned} \text{EMD}'(P, Q, D) &= \text{EMD}(P', Q', D') = (\text{from Corollary 1}) = \text{EMD}(\tilde{P}, \tilde{Q}, \tilde{D}) = \\ &= (\text{from definition (2.10) of } \text{EMD}^*) = \frac{\text{EMD}^*(P, Q, D)}{\max\{\sum P_i, \sum Q_j\}}. \end{aligned}$$

Thus, EMD^* is metric iff EMD' is metric. The latter's metricity, according to Theorem 1, requires equality of total masses of all network states and metricity of the ground distance. From the definition of EMD' , it is clear that *all* network states P' and Q' supplied to EMD by EMD' have the same total mass M . As to metricity of the ground distance, the identity of indiscernibles and symmetry straightforwardly follow from the corresponding properties of the original ground distance D and our choice of the ground distances to/from the bank nodes to be non-negative and symmetric. The triangle inequality trivially holds as network state extension does not introduce any new triangles. Hence, by Theorem 1, EMD' and, consequently, EMD^* is metric. ■

Having generalized Earth Mover's Distance, so that it can handle comparison of social network states, we will now be using EMD^* as our Earth Mover's Distance of choice in the definition (2.9) of Social Network Distance.

We have finished the design of SND and its components, and are now concerned with SND's efficient computation.

2.6 Efficient Computation of Social Network Distance

While we have designed SND (2.9), its computation is non-trivial. Since SND is a linear combination of multiple instances of EMD^* , its computation involves:

- Computing the ground distance $D(G(t), v)$ based on the structure of the underlying network $G = \langle V, E \rangle$ ($|V| = n$, $|E| = m$) and the opinions of the users in network state $G(t)$.
- Computing EMD^* , when the network states and the ground distance are provided.

Computing the ground distance D implies computing shortest paths in a network with edge weights $-\log \mathbb{P}_{ij}(P, v)$. Direct all-to-all shortest path computation using Dijkstra’s algorithm would incur time cost $O(n^2 \log n)$ for sparse G . Computing EMD^* is algorithmically equivalent to computing EMD , and, since the latter is formulated as a solution to a transportation problem, it can be computed either using a general-purpose linear solver, such as Karmarkar’s algorithm [74], or a solver that exploits the special structure of the transportation problem, such as the transportation simplex algorithm [75]. The time complexity of both these algorithms, however, is supercubic in n . Thus, the exact computation of SND using existing techniques is prohibitively expensive at the scale of real-world online social networks. Furthermore, the existing approximations of EMD are either inapplicable to the comparison of network states derived from a social network’s states, since they drastically simplify the ground distance [76, 77], or are effective only for some graphs, such as trees, structurally not characteristic of social networks [78]. Nevertheless, in what follows, we propose a method to compute SND exactly and in time linear in n under the following two realistic *assumptions*.

Assumption 1: The number n_Δ of users who change their opinions between two network states under comparison is significantly smaller than the total number n of users in the network.

Assumption 2: The log-likelihoods $-\log \mathbb{P}_{ij}(P, v)$ of opinion spread—being the edge weights in the network in which ground distances are defined as lengths of shortest paths—are positive

integers bounded from above by constant $U \ll +\infty \in \mathbb{Z}^+$.

Assumption 1 is reasonable in most applications the network states under comparison are not very far apart in time and, hence, $n_\Delta \ll n$; Assumption 2 is reasonable, because most of the log-likelihoods are large reals, and rounding them would not introduce a large error.

We will now use the above stated assumptions to design an efficient algorithm for SND. Since, according to its definition (2.9), SND's computation involves computation of four instances of EMD^* , we will actually be dealing with efficient computation of EMD^* on the inputs supplied by SND. Our method for efficient computation of SND requires the following two lemmas.

Lemma 1. *For any two network states $P \in \mathbb{R}^n$ and $Q \in \mathbb{R}^n$, and ground distance $D \in \mathbb{R}^{n \times n}$, if $P_i = Q_i = 0$, then the removal of i 'th elements from P , Q , as well as i 'th row and column from D does not affect the value of $\text{EMD}^*(P, Q, D)$.*

Lemma 1 is straightforward, since zero-value P_i and Q_i do not supply or consume any mass in the underlying transportation problem, and, hence, do not affect the cost of the optimal transportation plan. While Lemma 1 allows removing redundant suppliers and consumers from the underlying transportation problem, the following Lemma 2 allows to transform the network states, without affecting the value of EMD^* , exposing the redundant suppliers and consumers for removal.

Lemma 2 (Network State Reduction). *Given two arbitrary network states $P, Q \in \mathbb{R}^n$ and a ground distance $D \in \mathbb{R}^{n \times n}$, if D is semimetric⁴, then for any $i \in \{1, \dots, n\}$,*

$$\begin{aligned} \text{EMD}^*(P, Q, D) = \text{EMD}^*([P_1, \dots, P_{i-1}, P_i - \min\{P_i, Q_i\}, P_{i+1}, \dots, P_n], \\ [Q_1, \dots, Q_{i-1}, Q_i - \min\{P_i, Q_i\}, Q_{i+1}, \dots, Q_n], D). \end{aligned}$$

⁴A semimetric is a metric with the symmetry requirement dropped.

Proof. First, we will show that there is always an optimal plan f_{ij} in the problem of optimal mass transportation from \tilde{P} to \tilde{Q} over \tilde{D} such that $\forall 1 \leq i \leq n : f_{ii} = \min \{\tilde{P}_i, \tilde{Q}_i\} = M$, and, then, use such a plan to argue about the value of EMD^* .

Consider an arbitrary optimal transportation plan \hat{f}_{ij} , and assume that $\exists i \in [1; n] : \delta = M - \hat{f}_{ii} > 0$. We will now use \hat{f} to construct another optimal transportation plan f_{ij}^\dagger such that $f_{ii}^\dagger = M$. Initially, we put $f^\dagger = \hat{f}$ and, then, re-route mass flows in f^\dagger to eventually achieve the desired value of f_{ii}^\dagger .

Since, initially, $f_{ii}^\dagger < M$, the remaining at least δ units of mass should be distributed by \tilde{P}_i and consumed by \tilde{Q}_i to/from other consumers/suppliers. Among those, let us pick the ones that supply/consume the least amount of mass to \tilde{Q}_i and from \tilde{P}_i , respectively: $\ell = \arg \min_{j \neq i} f_{ji}^\dagger$, and $r = \arg \min_{j \neq i} f_{ij}^\dagger$. Without loss of generality, let us assume that $f_{\ell i}^\dagger \leq f_{ir}^\dagger$ and denote $\Delta = \min \{f_{\ell i}^\dagger, \delta\}$. Now, we will re-route Δ units of mass in f^\dagger as follows:

$$f_{\ell i}^\dagger \leftarrow f_{\ell i}^\dagger - \Delta, \quad f_{\ell r}^\dagger \leftarrow f_{\ell r}^\dagger + \Delta, \quad f_{ir}^\dagger \leftarrow f_{ir}^\dagger - \Delta, \quad f_{ii}^\dagger \leftarrow f_{ii}^\dagger + \Delta.$$

The above performed mass flow re-routing scheme is illustrated in Figure 2.7.

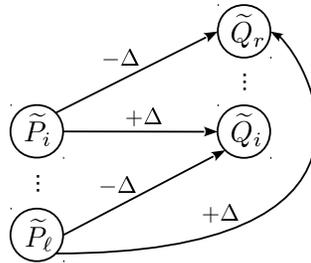


Figure 2.7: Re-routing of mass flow to convert an optimal transportation plan \hat{f} into plan f^\dagger .

The updated transportation plan is legal, as the total amount of mass supplied or consumed by

each node has not changed. The total cost of f^\dagger has been updated as follows

$$\begin{aligned} \text{newcost}(f^\dagger) &\leftarrow \text{cost}(f^\dagger) - \Delta\tilde{D}_{\ell_i} - \Delta\tilde{D}_{i_r} + \Delta\tilde{D}_{i_i} + \Delta\tilde{D}_{\ell_r} = \\ &= (\text{since } D \text{ and, hence, } \tilde{D} \text{ is semimetric, } \tilde{D}_{i_i} = 0) = \text{cost}(f^\dagger) - \Delta(\tilde{D}_{\ell_i} + \tilde{D}_{i_r} - \tilde{D}_{\ell_r}) \\ &\leq (\text{since } \tilde{D} \text{ is semimetric, } \tilde{D}_{\ell_i} + \tilde{D}_{i_r} \geq \tilde{D}_{\ell_r}) \leq \text{cost}(f^\dagger). \end{aligned}$$

Since the cost of the obtained legal plan f^\dagger cannot be strictly less than the cost of an optimal plan \hat{f} , the performed update of f^\dagger has not changed its cost, and the updated f^\dagger is still an optimal plan. The described above re-routing procedure is repeatedly performed on f^\dagger until f_{ii}^\dagger reaches $M = \min\{\tilde{P}_i, \tilde{Q}_i\}$.

Finally, to see why the statement of the lemma holds, we observe that the value of EMD^* is the cost of any optimal transportation plan, and the cost of f^\dagger in particular. However, the cost of f^\dagger does not depend on f_{ii}^\dagger , since, due to semimetricity of \tilde{D} , mass f_{ii}^\dagger gets transported at cost $\tilde{D}_{ii} = 0$. Thus, M can be subtracted from \tilde{P}_i , \tilde{Q}_i , and f_{ii}^\dagger , without affecting the total cost of f^\dagger . The solution of the latter modified transportation problem, however, is exactly

$$\text{EMD}^*([P_1, \dots, P_{i-1}, P_i - M, P_{i+1}, \dots, P_n], [Q_1, \dots, Q_{i-1}, Q_i - M, Q_{i+1}, \dots, Q_n], D).$$

■

We will, now, state the main result for the efficient computation of SND as Theorem 5, whose constructive proof provides the algorithm for SND's computation.

Theorem 5. *Under Assumptions 1 and 2, SND between network states $P = [P_1, \dots, P_n]$ and $Q = [Q_1, \dots, Q_n]$ defined over network $G = \langle V, E \rangle$, ($|V| = n, |E| = m$) can be exactly computed in time*

$$T = O(n_\Delta(m + n\sqrt{\log U} + n_\Delta^2 \log(n_\Delta n U))).$$

Proof. Throughout this proof, we will use notation $P^+ = P^{(+1)}$ to denote a network state where users holding negative opinions are considered neutral, and $D(P, +) = D(P, +1)$ to denote the ground distance for the spread of opinion $+1$ through the network in state P , as well as the similar notation P^- and $D(P, -)$ for negative opinions.

We will focus on the efficient computation of the first summand $\text{EMD}^*(P^+, Q^+, D(P, +))$ in definition (2.9) of $\text{SND}(P, Q, D)$, as computation of three other summands is algorithmically equivalent and takes the same time.

For the analysis of the computation of $\text{EMD}^*(P^+, Q^+, D(P, +))$, let us assume, without loss of generality, that $\sum_{i=1}^n P_i^+ \geq \sum_{j=1}^n Q_j^+$. As per (2.10), $\text{EMD}^*(P^+, Q^+, D(P, +))$ is the solution of a transportation problem with suppliers $\tilde{P}^+ = [P_1^+, \dots, P_n^+, \mathbb{0}_{1 \times n}]$, consumers $\tilde{Q}^+ = [Q_1^+, \dots, Q_n^+, Q^{+(1)}, \dots, Q^{+(n)}]$, and ground distance $\tilde{D}(P, +)$.

Now, we can apply Lemmas 1 and 2 to reduce the size of the obtained transportation problem. From Assumption 2, $\tilde{D}(P, +)$ is semimetric. Non-negativity and identity of indiscernibles straightforwardly follow from Assumption 2 and the definition of the length of a shortest path. Subadditivity follows from the shortest path problem's optimal substructure. Thus, we can apply Lemma 2 to each pair $\tilde{P}_i^+, \tilde{Q}_i^+$ of corresponding suppliers and consumers, and due to Assumption 1, a large number $(n - n_\Delta)$ of them have equal values. As a result, many suppliers and consumers become empty. Then, due to Lemma 1, all the obtained empty nodes can be disregarded. If we put $M_i = \min\{P_i^+, Q_i^+\}$, then the reduced transportation problem is defined for suppliers $[P_{i_1}^+ - M_{i_1}, \dots, P_{i_{n_\Delta}}^+ - M_{i_{n_\Delta}}]$ and consumers $[Q_{j_1}^+ - M_{j_1}, \dots, Q_{j_{n_\Delta}}^+ - M_{j_{n_\Delta}}, Q^{+(1)}, \dots, Q^{+(n)}]$, and ground distance $\tilde{D}(P, +)$ that contains only the rows and columns corresponding to the remaining suppliers and consumers. The remaining suppliers and non-bank consumers correspond to the users who have different opinion in P^+ and Q^+ , and the number of such users, due to Assumption 1, is at most n_Δ . The bank nodes $Q^{+(1)}, \dots, Q^{+(n)}$,

however, do not get affected by Lemma 2 in \tilde{Q}^+ (since only the banks of the lighter network state P can have non-zero mass) and hence are not removed, yet, they are removed from \tilde{P}^+ due to Lemma 1. Thus, we have an unbalanced transportation problem, where the number n_Δ of suppliers is much less than the number $n + n_\Delta$ of consumers.

Now, in order to compute $\text{EMD}^*(P^+, Q^+, D(P, +))$, we need to compute $\tilde{D}(P, +)$ and to actually solve the obtained transportation problem.

Due to the structure of the reduced transportation problem, we need to compute only a small part of $\tilde{D}(P, +)$. Since there are at most n_Δ suppliers, we need to solve at most n_Δ instances of single-source shortest path problem (SSSP) with at most $n_\Delta + n$ destinations. Since, due to Assumption 2, edge costs in the network are integer and bounded by U , each SSSP instance can be solved using Dijkstra's algorithm based on a combination of a radix and a Fibonacci heaps [30] in time

$$T_{sssp} = O(m + n \log \sqrt{U}).$$

(Notice, that if we assumed $\sum_{i=1}^n P_i^+ \leq \sum_{j=1}^n Q_j^+$, and the reduced \tilde{P}^+ contained $n_\Delta + n$ nodes, we would *not* need to run $n_\Delta + n$ SSSP instances. Instead, we would invert the edges in the network and compute the shortest paths in reverse, still solving only n_Δ SSSP instances.)

Next, we approach the solution of the reduced transportation problem with known ground distances. This problem can be viewed as a minimum-cost network flow problem in an unbalanced bipartite graph, where the number of consumers is much greater than the number of suppliers or vice versa. Since, due to Assumption 2, edge costs are integers bounded by U , our minimum-cost flow problem can be solved using Goldberg-Tarjan's algorithm [79] augmented with the two-edge push rule of Ahuja et al. [80] in time

$$T_{transp} = O(n_\Delta m + n_\Delta^3 \log(n_\Delta \max_{i,j} \tilde{D}(P, +)_{ij})).$$

Since no shortest path has more than $(n - 1)$ edge, and the edge costs are bounded by U , the expression for time simplifies to

$$T_{transp} = O(n_{\Delta}m + n_{\Delta}^3 \log(n_{\Delta}nU)).$$

Thus, the total time for computing $\text{EMD}(P^+, Q^+, D(P, +))$ and, consequently, $\text{SND}(P, Q, D)$ is

$$T = O(n_{\Delta}T_{SSSP} + T_{transp}) = O(n_{\Delta}(m + n \log \sqrt{U} + n_{\Delta}^2 \log(n_{\Delta}nU))).$$

■

Theorem 5 immediately entails the following corollary.

Corollary 2. *If the social network is sparse, that is $m = O(n)$, and the number n_{Δ} of users who changed their opinions is bounded, then SND is computable in time $O(n)$.*

2.7 Experimental Results

In this section, we report experimental results, demonstrating the utility of SND in applications in comparison to other distance measures. We also evaluate scalability of our implementation of SND .

2.7.1 Experimental Setup

Twitter Data: Our Twitter dataset is based on the crawled data of [81], and includes 48M tweets sent over 6 years. From these tweets, we select those sent between May-2008 and August-2011, containing hashtags related to the political topics—such as “Obama”, “GOP”, “Palin”, “Romney”—popular in the US during these years, and connect users in a network based on their follower-followee relationship. As a result, we obtain a network of 10k users

tweeting about politics, each having an average of 130 neighbors in the network. Within each quarter, we quantify the sentiment of each tweet as described in [82] using the sentiment classification approach of [51]. Then, we find the users who posted at least as many as 10% of the average number of per-user tweets posted within the quarter, and label them as active. Other users are assumed to be neutral, and their opinions for the quarter are set to 0. An active user’s opinion is set to +1 (−1) if he or she has posted at least 4 times more positive (negative) than negative (positive) tweets within the quarter, assuming that such a skew in the tweets’ sentiment is enough to identify whether the user likes or dislikes the topic. Otherwise—if an active user has posted enough of both positive and negative tweets—this user’s opinion is set to 0, that is, such user is considered neutral. As soon as we have quantified the quarterly opinion of each user, the opinions of all the users comprise that quarter’s network state.

Synthetic Data: We also perform experiments on synthetic scale-free networks of sizes $|V|$ from 10k to 200k and scale-free exponents from -2.9 to -2.1 . To generate the first network state, a number of initial adopters are chosen uniformly at random, and approximately equal numbers of them adopt opinions +1 and −1. Each subsequent network state $G(t+1)$ is randomly generated from the preceding network state $G(t)$ as follows. A number of $G(t)$ ’s neutral users get a chance to be activated. Each of them adopts an opinion from her neighbors with probability \mathbb{P}_{nbr} and a random opinion with a smaller probability \mathbb{P}_{ext} . If a user is to adopt an opinion from the neighbors, which opinion to adopt is decided in a probabilistic voting fashion based on the numbers of active in-neighbors of each kind. This generative model is a version of Independent Cascade model [52], where edges in a neighborhood are activated simultaneously with probability \mathbb{P}_{nbr} , and external influence \mathbb{P}_{ext} is allowed.

Distance Measures: In our experiments, SND does not make any assumptions regarding the above described generative process of opinion evolution in synthetic data, and assumes the simple model (2.1), whose parameter $c_{adverse}$, $c_{neutral}$, and $c_{friendly}$ values are learned from

how well SND performs in applications, and, as a result, are set, respectively, to 1000, 40, and 5 for anomaly detection experiments, and to 100, 20, and 5 in opinion prediction experiments. SND is compared with the following distance measures.

- *hamming*(P, Q). Hamming distance is a representative of ℓ_p -like distance measures performing basic coordinate-wise comparison.
- *quad-form*(P, Q, L) = $\sqrt{(P - Q)L(P - Q)^\top}$. Quadratic-Form Distance [57] based on the Laplacian matrix L [83] of the network. It combines the differences of opinions of the corresponding users based on the network's structure.
- *walk-dist*(P, Q) = $\frac{1}{n} \|con(P) - con(Q)\|_1$. Compares vectors $con(P) = [con(P_1), \dots, con(P_n)]$ of users' "contention", where $con(P_i)$ is the amount by which the i 'th user's opinion deviates from the opinion of this user's average active in-neighbor. Thus, *walk-dist* summarizes how different the network's users are from their respective neighbors.

2.7.2 Detecting Anomalous Network States

Synthetic Data: In a series $G(1), \dots, G(t), \dots$ of network states, we want to detect which transitions in the series are anomalous, that is, when opinions change unexpectedly deviating from their established evolution pattern. In particular, we are interested in those anomalies that are hard to detect by observing simple summaries of social network states, such as the number of newly activated users. To simulate such anomalies with synthetic data, we change the values of \mathbb{P}_{nbr} and \mathbb{P}_{ext} —controlling the process of opinion evolution from one network state to the other—preserving their sum, thereby, affecting which users get activated, yet, maintaining the same activation rate.

To detect anomalies, in a series of network states, we compute the distances between adjacent states, normalize these distances by the number of active users, and rescale the obtained

values to fit range $[0, 1]$. Then, spikes in the resulting distance series are considered anomalies.

A qualitative analysis of anomaly detection on synthetic data is presented in Figure 2.8.

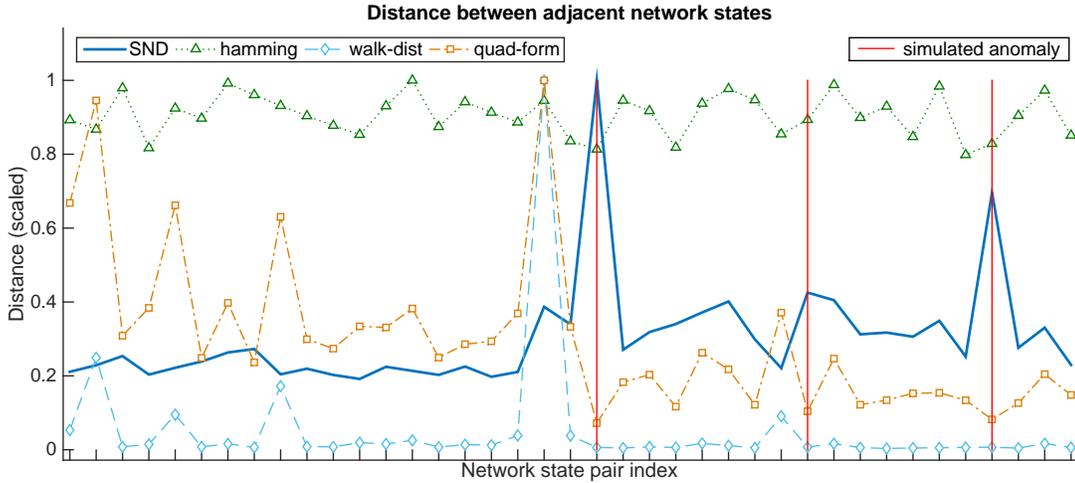


Figure 2.8: Anomaly detection on synthetic data. $|V| = 20k$, scale-free exponent $\gamma = -2.3$. A series of 40 network states is generated using $\mathbb{P}_{nbr} = 0.12$ and $\mathbb{P}_{ext} = 0.01$ for normal and $\mathbb{P}_{nbr} = 0.08$ and $\mathbb{P}_{ext} = 0.05$ for anomalous network states’ generation, respectively. The three simulated anomalies are displayed as solid vertical lines.

For each simulated anomaly, SND produces a well noticeable spike, reacting to the qualitative change in the opinion dynamics process, while other distance measures do not recognize such anomalies. (The additional experiment exposing this difference in sensitivity of SND vs. simpler distance measures is provided in Section 2.7.4.)

In order to quantify the performance of the competing distance measures at detecting simulated anomalies, we create a simple anomaly score $S_t = |(d_t - d_{t-1}) + (d_t - d_{t+1})|$, where d_t is the value of a given distance measure at time t normalized by the number of users active at time t and rescaled. We rank the network state transitions for each compared distance measure by S_t in decreasing order and compute true and false positive rates for increasing ranks. The corresponding ROC curves are displayed in Figure 2.9. SND’s accuracy dominates that of competing distance measures throughout the spectrum of false positive rates. Particularly, for false positive rates up to 0.3, SND achieves a true positive rate of 0.83, while the next best distance measure (*hamming*) achieves only 0.4.

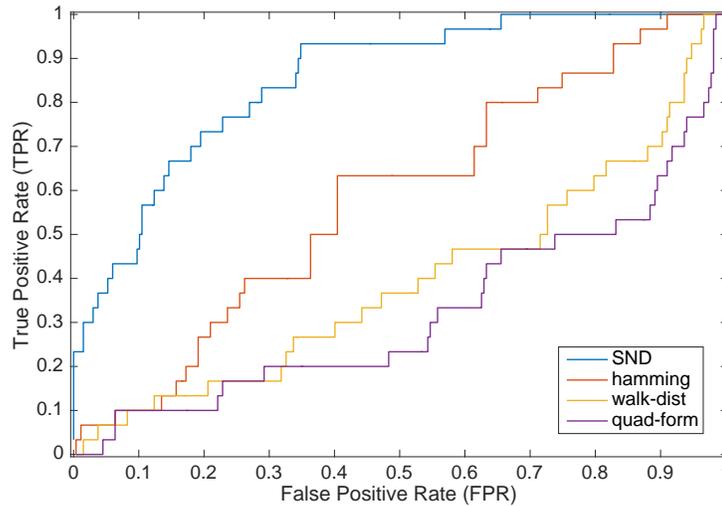


Figure 2.9: ROC curves comparing the quality of anomaly detection by different distance measures in a series of 300 network states over synthetic network with $|V| = 30k$ and scale-free exponent $\gamma = -2.3$. The network states are generated using $\mathbb{P}_{nbr} = 0.08$ and $\mathbb{P}_{ext} = 0.001$ for normal and $\mathbb{P}_{nbr} = 0.07$ and $\mathbb{P}_{ext} = 0.011$ for anomalous instances.

Twitter Data: To obtain the ground truth for anomaly detection on our Twitter dataset, we collect “search interest” data from Google Trends⁵ and cross-check this data with American Presidents⁶ log of political events in the US. The anomaly detection results for topic “Obama” are shown in Figure 2.10.

We can distinguish two types of events based on SND’s behavior relatively to that of other distance measures. One type is the polarizing events when SND noticeably disagrees with the other distance measures. For example, during quarters 05’09-11’09, the Economic Stimulus Bill had a highly polarized response in the House of Representatives⁷, with no Republican voting in its favor. Another such anomaly takes place during quarters 02’10-08’10, when the Affordable Care Act (“Obama Care”) was introduced, and which was and still remains a very controversial topic based on the House vote distribution⁸ and the analysis from socialmen-
tion.com⁹.

⁵<http://www.google.com/trends/explore>

⁶<http://www.american-presidents-history.com>

⁷<http://www.nytimes.com/2009/01/29/us/politics/29obama.html>

⁸<http://www.healthreformvotes.org/congress/roll-call-votes/h165-111.2010>

⁹<http://socialmention.com/search?q=obama+care>

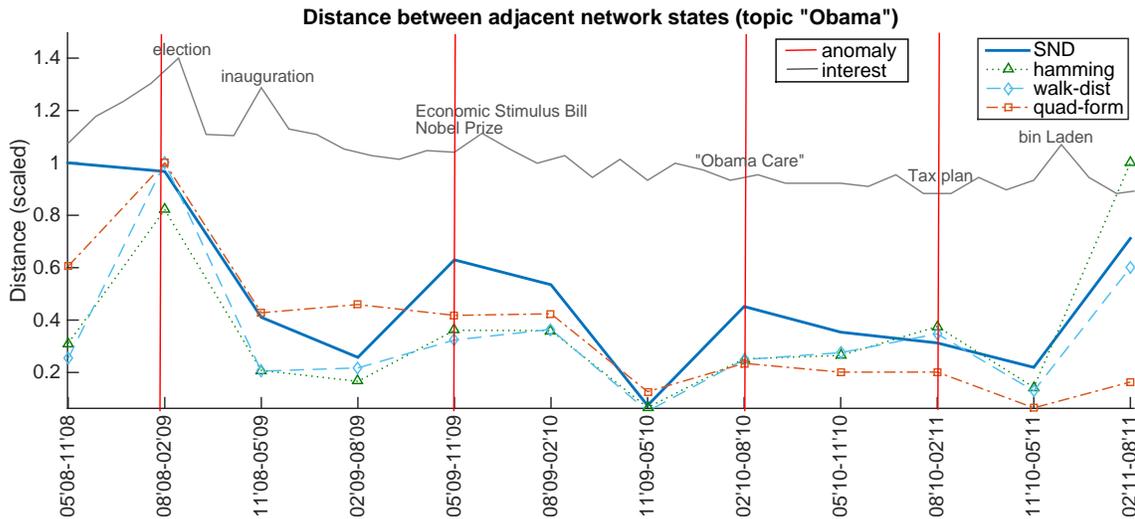


Figure 2.10: Anomaly detection on Twitter data (May’08-Aug’11). The distance series are accompanied by the curve showing Google Trends’ scaled interest in topic “Obama”. Network states detected to be anomalous by at least one distance measure are displayed as solid vertical lines.

The other events are those where SND agrees with the other distance measures. Three examples are (a) “election”, (b) “Tax plan”, and (c) “bin Laden” (even though, all distance measures noticeably increase their value during the last quarter, we do not mark this quarter as anomalous, since we do not have the distance values for the next quarter.)

The (a) election of Barack Obama as the President of the US, extensively covered by the news media, had likely been accompanied by a very noticeable change in the rate of new user activation, so, as expected, both SND and simpler distance measures sensitive to the user activation rate successfully detect this anomaly. However, the (b) Obama’s tax cut extension and (c) bin Laden’s assassination—not flagged as anomalies by SND—were not polarizing, as the tax cut had received large support in the Senate from both Democrats and Republicans¹⁰, while bin Laden’s assassination has probably evoked the same type of sentiment on Twitter across the US.

¹⁰<http://tinyurl.com/wiki-obama-tax-relief-2010>

$d_{-t} = \text{dist}(G(-t-1), G(-t))$ between adjacent past network states, then, extrapolate the obtained series of distances via linear 4-point¹¹ least squares fitting to estimate the expected distance d^* from the most recent $G(-1)$ to the yet not fully known true current network state $\widehat{G}(0)$. Then, we search for the assignment of opinions to the target users in the partially known current network state $G(0)$ —resulting in a candidate network state $G^*(0)$ —that would make the distance $\text{dist}(G(-1), G^*(0))$ from the most recent to the candidate network state as close to the estimate d^* as possible.

While there may be multiple candidate network states $G^*(0)$ whose distances are close to d^* —the network states in Figure 2.11 close to the sphere of radius d^* centered at $G(-1)$ —due to the spatially-sensitive nature of network state comparison that SND provides, the set of such candidate network states will be rather small, and the opinions of the target users in these network states will be close to the true target user opinions in $\widehat{G}(0)$.

We have used two methods for the search of the best opinion assignment. The first method is a randomized search with uniformly randomly chosen opinions, and the number of random opinion assignments (100 in our experiments) being considerably lower than the total number of possible assignments (1M+ in our experiments). The second method was greedy hill climbing, the results for which are not reported, performing no better than the randomized search.

In each experiment, we uniformly randomly select 20 active users—with roughly equal representation of positive and negative opinions—in the current network state, predict their opinions and measure the prediction accuracy. We repeat this procedure 10 times, each time targeting a different set of users, and report means and standard deviations of the obtained prediction accuracies.

The predictions are made using the above distance measure-based method with SND as well as other distance measures. To put the prediction performance of these methods in context,

¹¹We have experimented with different numbers of points used with least squares. Using less than 4 points resulted in a poor opinion prediction performance, while using more than 4 points did not improve the performance much.

we include in the comparison several non-distance measure-based opinion prediction methods.

- *icc-simulation*, *ltc-simulation* [84] simulate the model—Independent Cascade or Linear Threshold with uniformly randomly chosen thresholds, respectively—until convergence multiple times (from 10 to 500, in our experiments) and use the modes of the target users’ opinion as the prediction. In our experiments, the simulation starts with the most recent completely known network state G_{-1} , and proceeds until 99.99% of users get active. The edge activation probabilities \mathbb{P}_{edge} are selected uniformly, with \mathbb{P}_{edge} ranging from 0.001 to 0.01. The results for the best \mathbb{P}_{edge} are reported.
- *icc-max-likelihood*, *ltc-max-likelihood* are max-likelihood-based methods similar to the method of Saito et al. [85]. Like SND-based opinion prediction, this method generates uniformly random opinion assignments, computes the likelihood of each resulting network state and uses the most likely one for the prediction. The opinion adoption likelihoods are computed as described in Section 2.2.2, with all edges assumed to be active, and $\varepsilon = 0.01$.
- *community-lp* [86, IV.B] detects communities in the network via label propagation and, then, predicts user opinions based on these users’ membership in the discovered communities. This method does not rely on any opinion dynamics model, and only assumes that users likely connect with other likeminded users.

We experiment with both synthetic and Twitter data. For synthetic data, we generate a scale-free network with $n = 10,000$ users and scale-free exponent $\gamma = -2.5$. A series of network states is generated using the same version of Independent Cascade model as was used in anomaly detection experiments, with probabilities of opinion adoption from the neighbors \mathbb{P}_{nbr} and from the external source \mathbb{P}_{ext} ranging between 0.001 and 0.2. The number of initially active users is set to 800.

User Opinion Prediction Accuracy, %				
Method	Synthetic Data		Twitter Data	
	μ	σ	μ	σ
SND	74.33	2.65	75.63	5.60
hamming	68.44	12.34	68.13	5.80
quad-form	66.67	13.58	67.50	9.63
walk-dist	56.22	15.35	31.88	9.98
icc-simulation[84]	76.25	9.54	59.38	4.17
ltc-simulation[84]	67.50	11.65	58.75	5.18
icc-max-likelihood[85]	67.41	7.03	57.50	8.02
ltc-max-likelihood[85]	57.50	8.45	55.63	11.78
community-lp[86]	65.25	9.43	56.87	8.43

Table 2.2: Means μ and standard deviations σ of user opinion prediction accuracies.

The opinion prediction results are summarized in Table 2.2. There are four important observations:

(a) Among the distance-based methods, SND always performs best on both synthetic and Twitter data, with the mean prediction accuracy of 74-75% and a consistently low standard deviation. This suggests that SND captures more opinion dynamics-specific information than other distance measures, and should be preferred, particularly, when such simple statistics as the rate of new user activation are uninformative.

(b) Among the non-distance measure-based methods, *icc-simulation*'s prediction accuracy on synthetic data is 76.25%, the best result, comparable to 74.33% accuracy of SND. Such good performance of *icc-simulation* on synthetic data is unsurprising, as its predictions rely on the Independent Cascade model, a version of which produced the synthetic data. On Twitter data, however, *icc-simulation*'s accuracy is 59.38%, compared to 75.63% accuracy of SND.

(c) Methods *icc-max-likelihood* and *ltc-max-likelihood* perform worse than SND, as they base prediction on the opinions of the closest active neighbors, while SND is looking for the most likely opinion propagation through potentially long paths in the network.

(d) SND-based method outperforms *community-lp* based on community detection via label propagation. In our experiments, *community-lp*'s prediction accuracy is 57-65%, while this method's authors report the accuracy of 95% for their data [86]. The likely cause of such a discrepancy is a *very* high level of homophily in their data (the reasons of which were discussed in [87]), while in our less homophilous data, *community-lp* performs worse by capturing only users' reachability by the opinions of each kind, whereas SND performs better by looking for the *most likely* opinion propagation scenario.

2.7.4 Sensitivity to Opinion Dynamics Models

In the anomaly detection and user opinion prediction experiments, performance of SND stemmed from its being spatially-sensitive to the changes in the user opinion distribution, and promptly reacting to qualitative changes in the underlying opinion spread process. In this section, we conduct an experiment confirming that sensitivity of SND. We show the effectiveness of SND in detecting qualitative changes in the user opinions' evolution under an advanced opinion dynamics model, that cannot be spotted by the distance measures performing coordinate-wise comparison. To that end, we generate a number of pairs $\langle G(1), G(2) \rangle$ of network states adjacent in time ($G(2)$ is generated from $G(1)$) over a synthetic scale-free network. Some of these pairs correspond to *normal transitions*, while others correspond to *anomalous transitions* in the network's evolution. For the normal transitions, $G(2)$ is generated from $G(1)$ using the Independent Cascade model [52]. For the anomalous transitions, most new user activations in $G(2)$ happen randomly, independently of the network's structure. We study the distances assigned to normal and anomalous network state transitions by SND and ℓ_1 , and plot them as functions of the number n_Δ of users whose opinions change over each network state transition. The results are shown in Figure 2.12.

We see that SND clearly separates anomalous transitions from normal ones, while ℓ_1 cannot

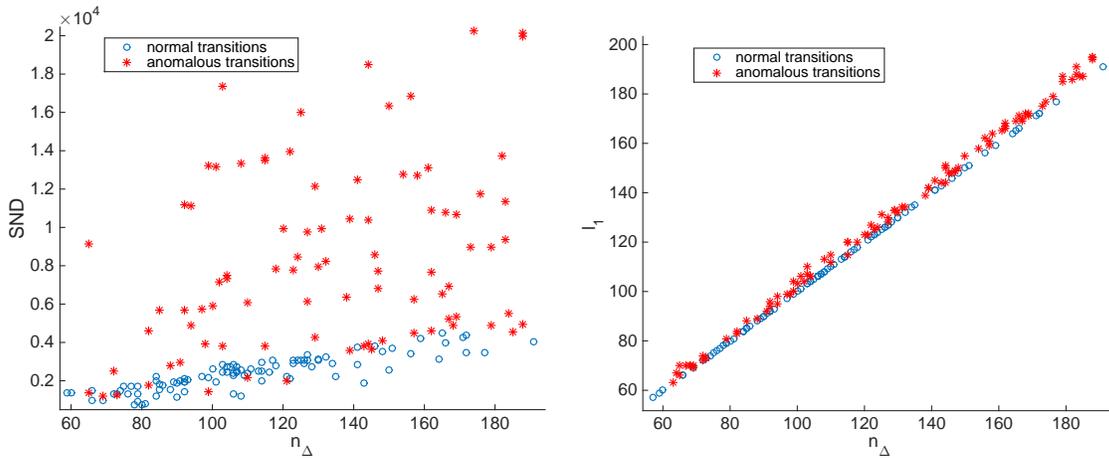


Figure 2.12: SND and ℓ_1 distances between network states of normal (ICC) and anomalous (random) transitions.

discern anomalous network state transitions, as ℓ_1 's value is mostly determined by n_Δ , which is representative of the distance measures performing coordinate-wise comparison.

2.7.5 Scalability of Social Network Distance

We have implemented¹² SND in MATLAB and C++. We use the minimum-cost network flow solver CS2 [88] that implements Goldberg-Tarjan's algorithm [79], but, unlike it is prescribed by Theorem 5, does not use the two-edge push rule of Ahuja et al. [80]. Additionally, for computing shortest paths, our implementation of Dijkstra's algorithm uses a priority queue based on a binary heap, rather than a combination of a Fibonacci and a radix heaps [30]. As a result, our implementation of SND scales slightly worse than linearly—as guaranteed by Theorem 5—but still very well to be applicable to real-world social networks. Figure 2.13 shows how our implementation of SND scales with respect to the number n of users in the network in comparison to a direct computation of SND using CPLEX' linear solver [89]. Our implementation's scalability with respect to the number n_Δ of users holding different opinions in two network states under comparison is shown in Figure 2.14.

¹²<http://cs.ucsb.edu/~victor/pub/snd/>

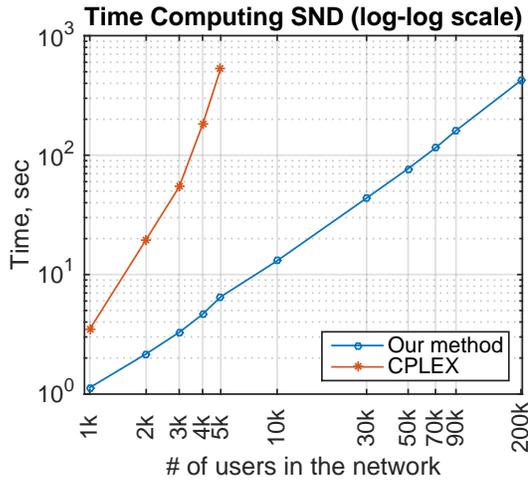


Figure 2.13: Time for computing SND when the number of users having different opinion is fixed at $n_{\Delta} = 1000$ and the total number of users n in the network grows up to 200k.

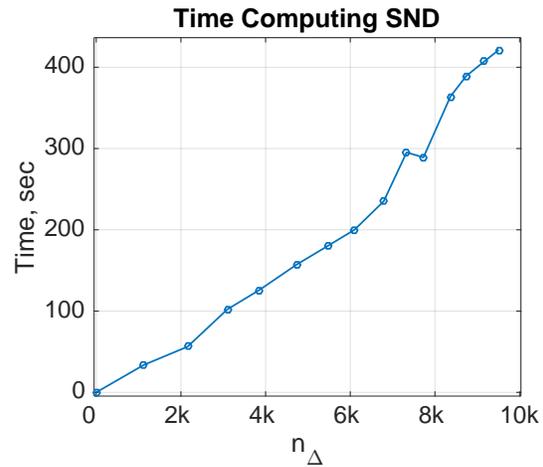


Figure 2.14: Time for computing SND using our method, with the network size fixed at $n = 20k$, and the number n_{Δ} of users who changed their opinions growing up to 10k.

2.8 Limitations

Despite the demonstrated effectiveness and efficiency of SND, there are scenarios in which its use is either prohibitively or unnecessarily expensive.

- *When SND is Overly Expensive:* One reason to choose a simpler distance measure, such as ℓ_p , over SND is the latter's computational cost. While it is asymptotically linear in the number n of nodes, its cost can potentially be too high in practice for networks having 100M+ nodes. In such networks, a single computation of SND can take several days. If it is nonetheless desirable to use SND on such a large network, one can partition that network into clusters of tractable size and perform the SND-based analysis on each individual cluster.
- *When SND is Superfluous:* Using SND may be superfluous if the changes in the rate of new user activation reveal enough information for the target application. For example,

the user activation rate alone is clearly enough to detect a presidential election day. For detection of such “anomalies”, a distance measure as simple as Hamming distance may suffice.

2.9 Conclusion

In this chapter, we focused on the design of a method for the analysis of the observed process of polar opinion dynamics in a social network. More specifically, we proposed Social Network Distance (SND)—the first distance measure for comparing the states of a social network containing polar opinions. Our distance measure quantifies how likely one state of a social network has evolved into another state under a given model of opinion dynamics. Despite the high computational complexity of the transportation problem underlying SND, we propose a linear-time algorithm for its exact computation, making SND applicable to real-world online social networks. We demonstrate the usefulness of SND in detecting anomalous network states and predicting user opinions, where SND-based methods consistently outperform competitors. Our anomaly detection method achieves a true positive rate (TPR) of 0.83 when the false positive rate (FPR) is 0.3, while the next best method’s TPR is only 0.4 at the same FPR. The accuracy of SND-based method for user opinion prediction in Twitter data is 75.63%, which is 7.5% higher than that of the next best method. We also show that, unlike the distance measures performing coordinate-wise comparison, SND can detect qualitative changes in the network’s evolution pattern. SND is a powerful alternative to simpler distance measures, and is effective when such summaries of network users’ behavior as the number of active users are uninformative, and a deeper insight into the opinion dynamics process is required.

2.10 Future Research

Among the directions for future research are the following.

- *New applications:* Since SND is, effectively, the first distance measure designed specifically for the comparison of states of a social network containing competing opinions, one potential future research direction is using SND in other applications operating in a metric space setting, such as network state classification, clustering, and search.
- *Combining macro-level distance measure-based and user-level analysis:* It may be lucrative to combine SND with non-distance measure-based methods. For example, in the method of Conover et al. [86] that predicts opinions based on the content of the users' tweets, the objective function can be augmented with an SND-based term, thereby, performing opinion fitting at both the micro-level of each user and the macro-level of the entire network.
- *Design of a distance measure for both structural and opinion changes:* Finally, it may be fruitful to design a distance measure that would capture changes in both the opinions of the users and the structure of the social network simultaneously. Such a distance measure would be more computationally complex than SND, but would result in a more accurate analysis when the network structure changes a lot from network state to network state.

Chapter 3

Non-Linear Models for Polar Opinion Dynamics in Social Networks

For decades, scientists have studied opinion formation in social networks, where information travels via word of mouth, with the particularly interesting case's being that of polar opinions and their competition in the network. The central problem addressed in this chapter is to design and analyze a general model that would capture how polar opinions evolve in the real world.

We propose a *general non-linear model of polar opinion dynamics*, rooted in several theories of sociology and social psychology. The model's key distinguishing trait is that, unlike in the existing linear models, such as DeGroot and Friedkin-Johnsen models, *an individual's susceptibility to persuasion is a function of his or her current opinion*. For example, a person holding a neutral opinion may be rather malleable, while "extremists" may be strongly committed to their current beliefs. The same model can be seen as a model of non-linear heat flow in a network-structured material. We also study three specializations of our general model, whose susceptibility functions correspond to different socio-psychological theories.

We provide a comprehensive theoretical analysis of our non-linear models' behavior using several tools from non-smooth analysis of dynamical systems. To study convergence, we use

non-smooth max-min Lyapunov functions together with the generalized Invariance Principle. For our general model, we derive a sufficient condition for the convergence to consensus. For the specialized models, we provide a full theoretical analysis of their convergence—whether to consensus or disagreement. The obtained results are rather general and easily apply to the analysis of other non-linear models defined over directed networks, with Lyapunov functions constructed out of convex components.

This chapter is organized as follows. Having motivated our work in Section 3.1, we review existing related opinion dynamics models in Section 3.2. Then, in Section 3.3, we define our models of polar opinion dynamics. Section 3.4 provides preliminaries from non-smooth analysis necessary for the convergence analysis of our models in Section 3.5. We complement our theoretical analysis of the models' behavior with simulation results, provided in Section 3.6. Finally, we conclude the chapter with a discussion of our results in Section 3.7.

3.1 Introduction

Studying the evolution of opinions of a group of people—the agents—connected in a directed social network, we assume that the objective means for opinion evaluation are limited, and the agents evaluate their opinions by comparison with the opinions of others [90]. Thus, the process of opinion formation in a group is a *network process*, where each agent's opinion changes due to the agent's interaction with his or her neighbors in the network.

In particular, we focus on *polar opinions*, which describe either degrees of proclivity toward one of two competing alternatives (e.g., Democrats vs. Republicans or iOS vs. Android) or an attitude—from extreme unfavorable to neutral to extreme favorable—toward a single issue (e.g., using nuclear power as an energy source). We will use the terms *opinion* and *attitude* interchangeably, and refer to them both as *an agent's state*. Our emphasis on polar opinions will manifest itself in that the agents in our non-linear models will change their opinion-adoption

behavior as their opinions shift toward one or another pole of the opinion spectrum.

In what follows, we will review the main components of an opinion dynamics model and connect them with the existing theories of sociology and social psychology, preparing a foundation upon which our models will be constructed.

Opinion formation via weighted averaging: The most basic network model of opinion dynamics is the weighted averaging model of DeGroot [91] (whose continuous-time version was studied earlier by Abelson [49]):

$$x(t+1) = Wx(t),$$

where t is time, $x(t)$ is a vector of agent states, and W is the row-stochastic adjacency matrix of the social network, with W_{ij} indicating the relative extent to which agent j influences the opinion of agent i , or, alternatively, the relative share of $x_j(t)$ in $x_i(t+1)$. According to this model, each agent forms his or her opinion as a weighted average of all the opinions available in the agent's out-neighborhood in the network.

The appeal of DeGroot model stems from its consistency with such theories of social psychology as social comparison theory [90], cognitive dissonance theory [92], and balance theory [93, 94], whose unifying idea is that the agents act to achieve balance with other group members or, alternatively, to relieve psychological discomfort from their disagreement with others. However, one limitation of DeGroot model is that the agents' "behavior" does not change depending on the agent, its current state—opinion or attitude—and the issue at hand.

Models with susceptibility to persuasion: At the very least, the strength of an agent's attachment to his or her opinion depends on the extent to which the issue is important to that agent and is representative of his or her values. Such a dimension of the strength of attitude—a

function of the agent and the issue—has arisen in multiple studies under the names of embeddedness [95], ego preoccupation and ego involvement [96, 97], among others. Friedkin-Johnsen model [98] addresses the limitation of DeGroot model by allowing the agents to have different *susceptibilities to persuasion*:

$$x(t + 1) = AWx(t) + (I - A)x(0),$$

where t , $x(t)$ and W are defined as before, I is the identity matrix, and A is a constant diagonal matrix whose diagonal element A_{ii} describes the extent to which agent i 's opinion is affected by the opinions of other agents as opposed to his or her own initial opinion. The diagonal elements of matrix $(I - A)$ are usually referred to as the agents' degrees of *stubbornness*. Friedkin-Johnsen model improves upon DeGroot model not only in terms of the model's interpretation, but also in terms of the model's behavior—while the typical asymptotic behavior of DeGroot model in a “well-connected” social network is the convergence to consensus, in case of Friedkin-Johnsen model, agents usually disagree. The latter behavior usually occurs in the real world [99].

State-dependent susceptibility in linear models: The “state-oblivious” definition of the agents' constant susceptibilities to persuasion of Friedkin-Johnsen model does not capture another component of the agents' strength of conviction, known in social psychology as commitment [95, 96, 100] or certainty [97], which is determined by each agent's attitude toward the issue. The dependency of susceptibility to persuasion on the agents' beliefs has been studied in the context of Friedkin-Johnsen model [101], where the asymptotic behavior of the model was empirically evaluated under several definitions of susceptibility $A(x(0))$ as a function of the initial opinions $x(0)$ of the agents.

Among the existing theories of social psychology, there is no agreement upon a single

correct definition of susceptibility A based on the agents' beliefs. One factor that has arisen in multiple studies as closely related to the strength of conviction is the attitude extremity or polarity [102, 103, 95, 104, 105, 106]. The conclusion that can be drawn from these works is that extreme opinions are more resistant to change, possibly, due to the preferential evaluation of attitude-congruent information by the agents holding extreme opinions. Alternatively, social comparison theory [90] suggests that, when the majority of the agents hold a certain, say, neutral, opinion, establishing a social norm, then the agents with opinions close to that norm have relatively weaker tendencies to change their positions, while the extreme opinions are unstable.

Existence of multiple alternative theories regarding the factors determining the strength of the agents' commitment to their opinions is not surprising, particularly, because these factors have been shown to be domain-specific [95]. Hence, it is rational to either study the opinion formation process under the most general definition of the strength of the agents' attitudes or to use multiple definitions of the attitude strength based on the existing socio-psychological theories. In this chapter, we will study both the most general definition of agent susceptibility as well as several specialized definitions consistent with different socio-psychological theories.

State-dependent susceptibility in non-linear models: The definition of susceptibility A as a function of the *initial state* $x(0)$ is beneficial in that it does not take away the model's linearity and, hence, allows application of the existing linear-algebraic techniques to the formal analysis of the model's asymptotic behavior. However, definition of A as a function of the *current state* $x(t)$, while would make the model non-linear, has at least two advantages. For one thing, the definition $A = A(x(t))$ may be more appropriate when the evolution of opinions is studied at a large time scale, as in the case when a group of people is working on a year-long project, and hardly anyone remembers what their opinions were a year ago. For another thing, and more importantly, in several existing studies [107, 108, 109], the agents' attitudes

are posited to be “constructed on demand”, and, in particular, according to the potentiated recruitment model [107], the strength of attitude is an emergent property of the process of attitude construction occurring when the attitude is recruited. Thus, in our models, we adopt the definitions of agent susceptibility dependent on the agents’ current states.

Contributions: We propose novel non-linear models of polar opinion dynamics and formally analyze their behavior. Our specific contributions are as follows.

- *Novel Models:* We propose a general non-linear model of polar opinion dynamics, where the agents’ susceptibilities to persuasion are general functions of the agents’ current beliefs. Additionally, we propose three specialized instances of the general model, having different definitions of agent susceptibility $A(x(t))$ corresponding to different theories of social psychology. The proposed models are novel in that they capture more traits of the opinion formation process than the existing models, and manifest a behavior unobserved in their linear counterparts—we can generally observe either the agents’ convergence to consensus or their persistent disagreement, depending on the agents’ initial beliefs $x(0)$.
- *Analysis of the General Model:* For our general model of polar opinion dynamics we prove a contraction property of its trajectories, and provide a sufficient condition for the convergence to consensus. That sufficient condition is rather general and, roughly, states that convergence to consensus takes place if the agents non-responsive to persuasion have identical states. The latter entails that, quite naturally, a disagreement among the agents may arise only if there are multiple agents having different beliefs and unwilling to change them.
- *Analysis of the Specialized Models:* We provide a comprehensive theoretical analysis of the asymptotic behavior of our specialized models, characterizing all their states of equilibrium—corresponding to either the states of consensus or disagreement—through

the analysis of certain partitions of the network, and prove each model's convergence. For the cases when a model converges to a state of disagreement, we provide an explicit expression for that limiting state, which depends on the network's structure as well as the beliefs and locations of the agents non-responsive to persuasion, yet, does not depend on the initial beliefs of the susceptible agents.

- *Novel Analysis of Convergence:* The standard tools for the analysis of convergence of non-linear continuous-time models, such as Lyapunov's Second Method [110] and (Barabashin-Krasovskiyy-)LaSalle Invariance Principle [111, 112], require existence of a smooth Lyapunov function, which may not and, sometimes, provenly does not exist [113] for a model defined over a directed network. In this chapter, we use max-min non-smooth Lyapunov functions along with several existing non-smooth analysis techniques to prove convergence of our non-linear models. While such Lyapunov functions have appeared in existing literature, to the best of our knowledge, we are the first to provide a full formal analysis of convergence of a continuous-time non-linear system defined over a directed network using such functions together with the generalized Invariance Principle.

3.2 Related Work

The numerous existing opinion dynamics models can be roughly divided into two groups: *analytic* and *algorithmic*.

3.2.1 Analytic Models

Analytic models are represented as systems of difference or differential equations

$$x(t+1) = W(x(t), t)x(t) \quad \text{or} \quad dx/dt = \dot{x} = W(x(t), t)x$$

and describe a process of agent interaction usually targeting a certain form of agreement among the agents. These models mainly differ in the extra properties of the agent interaction process besides the agents' effort toward reaching agreement. All our models proposed in this chapter belong to this group.

Analytic models have long been studied by sociologists, starting with the early works of French [114] and Harary [115]. Nowadays, the basic formulation of the weighted averaging process is usually referred to as DeGroot model [91]. A variation of DeGroot model with some agents' states kept constant has been studied by Pirani and Sundaram [116]. An improvement upon DeGroot model was proposed by Friedkin and Johnsen [98], who enabled the agents to have individual levels of susceptibility to persuasion by other agents. Variations of Friedkin-Johnsen model have been studied in a context of a group's discussing a sequence of issues [117, 118]. The question of the dependence of A upon the agents' beliefs has been empirically studied by Friedkin in [101]. A variation of Friedkin-Johnsen model with time-dependent $A(t)$ and its connection to the underlying notion of dissonance minimization was discussed in work [119] by Groeber et al. A Friedkin-Johnsen-type model with stubborn agents has been studied as a local interaction game by Ghaderi et al. [120].

Another type of analytic models—close in spirit to our models—are the models with state-dependent agent interaction, $W(x)$, and, in particular, the bounded confidence models [121, 122], whose key idea is that only the agents with close enough states can interact. The two popular representatives of such models are Hegselmann-Krause (HK) [123] model and the model of Deffuant et al. [124]. Some convergence results for HK model have been proven by Blondel et al. [125] and MirTabatabaei and Bullo [126]. Sufficient convergence conditions for a more general model with state-dependent agent interaction, that includes HK model as a special case, have been studied by Lorenz [127].

A special subtype of bounded confidence models are those that allow for stubbornness, leadership, antagonism, or zealotry, and whose behavior has been investigated through sim-

ulation. In particular, in [128], Deffuant et al. study the behavior of Deffuant’s model with smooth confidence bounds in the presence of stubborn agents. Kurmyshev et al. [129] extend Deffuant’s model with two types of agents characterized by “friendly” and “antagonistic” interaction, respectively. Jalili [130] has studied the effect of the choice of the subset of stubborn leaders as well as a particular network structure upon the bounded confidence model’s convergence rate. Sobkowicz [131] has considered the “Deffuant model with emotions”, where different opinions have varying resilience to change. In particular, the author considered the cases when the extreme opinions are more resilient to change than the neutral opinions, as well as the case of an asymmetric dependency of the opinion resilience on the opinion value. These opinion resilience mechanisms are similar to some of those we use in our specialized models in Section 3.3. Chen et al. [132] investigated how stubborn leaders can attract followers in the context of a bounded confidence model that incorporates the leaders’ reputation, stubbornness, appeal, as well as the extremity of their opinions. Finally, Tucci et al. [133] have studied the bounded confidence model with stubborn leaders and investigated the effect of the number of leaders on the opinion dynamics profile.

Similar analytic models are studied in the control systems and robotics communities, in the context of multi-agent coordination problems [134, 135]. The models with time-varying topology $W(t)$ of the network have been studied in [136, 137, 113, 138]; the models for signed networks, allowing for agents’ friendly and antagonistic interaction, have been studied in [139, 140]; the models with randomized agent interaction have been studied in [141, 142].

The final class of analytic models are the models considered in the context of the naming game. Specifically, Waagen et al. [143] design a naming game model with discrete opinions and zealots—who do not change their opinion—and study the effect of the number of zealots on the opinion dynamics of the entire population.

3.2.2 Algorithmic Models

Algorithmic models for opinion dynamics are usually defined as probabilistic or deterministic combinatorial algorithms describing how the agents update their states. These models usually operate with discrete agent states and in discrete time. A notable difference of algorithmic models from their analytic counterparts is that algorithmic models are usually data-driven, that is, such models are usually to be fit to data, whereas the analytic models are “prescriptive”. One model in this group is the Independent Cascade Model [144], where the agents get “activated” with an opinion by their neighbors in a probabilistic fashion. The basic version of this model uses binary opinions—indicating presence or absence of an opinion—and is usually used in the context of the influence maximization problem [145]. A version of the Independent Cascade Model for the case of multiple competing opinions has been proposed in [52]; a version with asynchronous communications has been studied in [146]. A related, yet more general model, allowing for competing opinions, is the switching-selection model of [147].

Two other types of algorithmic models are the Voter model [148, 149, 150, 151] and the Linear Threshold model [152, 153], where in the former model, each agent is activated in a probabilistic fashion based on the number of active agents in the neighborhood, and in the latter model, agents become active as soon as the number of active neighbors surpasses a constant threshold. Versions of the Linear Threshold model for the case of competing opinions have been studied in [54]. The extensions of the discrete-opinion Voter model with stubborn agents have been considered in works [154, 155, 151], where the authors studied the models’ long-run behavior as well as the problem of influence maximization. Finally, there are Bayesian algorithmic models [156], whose agent state updates are based on the Bayes rule.

3.3 Models for Polar Opinion Dynamics

3.3.1 General Model

We define the *general model of polar opinion dynamics* as follows:

$$\dot{x} = -A(x)Lx. \quad (3.1)$$

where $x(t) \in [-1, 1]^n$ represents the agents' states, $A(x(t)) \in \text{diag}([0, 1]^n)$ is a diagonal matrix whose diagonal elements are the agents' state-dependent and possibly different susceptibility functions locally Lipschitz in $[-1, 1]^n$, $L = I - W$ is the network's Laplacian matrix, and $W \in [0, 1]^{n \times n}$ is the row-stochastic adjacency matrix of the directed network, with its edge weight W_{ij} measuring the amount of relative influence of agent j upon agent i .

The mathematical interpretation of the above defined model is as follows. In (3.1), the negative Laplacian $-L$, when applied to $x(t)$, measures how much, on (weighted) average, the agent's state is smaller than the states of the agents in its out-neighborhood $N_{out}(i) = \{j \mid j \neq i \wedge W_{ij} > 0\}$

$$(-Lx)_i = \sum_{j \in N_{out}(i)} w_{ij}(x_j - x_i).$$

When $(-Lx)_i > 0$ and agent i is *open* to persuasion, that is, $A_{ii}(x) > 0$, then $\dot{x}_i > 0$, and x_i grows, "following" its out-neighbors. Conversely, if $(-Lx)_i < 0$, the state of an open agent i decreases. If either an agent's state is in balance with the states of its out-neighbors, or the agent is *closed* to persuasion, that is, $A_{ii}(x) = 0$, then this agent's state does not change.

Model (3.1) can also be thought of as a non-linear generalization of the heat diffusion model $\dot{x} = \alpha \Delta x$, where the negative Laplacian $-L$ of (3.1) corresponds to the finite-difference approximation of the continuous-space Laplace operator Δ , and the rate $A(x)$ at which the state

of the model evolves may be thought of as the temperature-dependent thermal diffusivity—a naturally occurring phenomenon [157].

The general model for polar opinion dynamics $\dot{x} = -A(x)Lx$ consists of two conceptual components: the averaging component $-Lx$ drives the agents towards agreement, while the susceptibility component $A(x)$ impedes this convergence process. The averaging component is based on such theories of social psychology as social comparison theory [90], cognitive dissonance theory [92], and balance theory [93, 94], whose unifying idea is that the agents act to achieve balance with other group members. The general idea of agents' susceptibility or stubbornness to persuasion comes from the socio-psychological studies of the strength of attachment to one's opinion [95, 96, 97]. The dependency of the agents' susceptibility $A(x)$ on their current beliefs agrees with the socio-psychological studies [107, 108, 109], that posit that the agents' attitudes are "constructed on demand".

3.3.2 Specialized Models

In addition to the general model (3.1), we will consider three specialized models, each with a different definition of state-dependent susceptibility $A(x(t))$, having different socio-psychological interpretations.

(i) *Model with Stubborn Extremists*: The first specialized model draws from the socio-psychological studies [102, 103, 95, 104, 105, 106] of the attitude extremity as being a major factor defining the strength of conviction, and whose definition of agent susceptibility $A(x) = (I - \text{diag}(x)^2)$ assumes that extreme opinions are more resistant to change than neutral opinions:

$$\dot{x} = -(I - \text{diag}(x)^2)Lx. \quad (3.2)$$

This model is appropriate when the extreme opinions *compete* in that an agent's strong prefer-

ence of one extreme implies this agent’s likely rejection of the opposite extreme. For example, this may be the case when agents’ states describe the degrees of support for one of the two major political parties in the US—inveterate Republicans or Democrats are unlikely to change their political affiliation, while neutral voters can be successfully attracted toward one or another pole of the opinion spectrum.

(ii) *Model with Stubborn Positives*: The second specialized model is a variation of the model with stubborn extremists with the asymmetric susceptibility function $A(x) = \frac{1}{2}(I - \text{diag}(x))$, where the agents only at one end of the opinion spectrum are stubborn:

$$\dot{x} = -\frac{1}{2}(I - \text{diag}(x))Lx. \quad (3.3)$$

This definition—inspired by the “Stubborn Left” and “Stubborn Right” susceptibility functions of Friedkin [101]—fits those cases when the agents at one, say, negative extreme of the opinion spectrum have no reason to reject the alternative opinion, while the agents having the opposite, positive, opinion have an incentive to maintain their position. For example, the opinion may describe the degree of liking for one of two smartphone brands, where opinion -1 corresponds to the neutrally marketed brand, while opinion $+1$ corresponds to the brand that is aggressively marketed not just as the best, but also as the only viable option.

(iii) *Model with Stubborn Neutrals*: Finally, in our third specialized model, drawing from the social comparison theory [90] and the studies of social norms [101], we defined agent susceptibility as $A(x) = \text{diag}(x)^2$, assuming that the neutral opinions are resilient to change, while the extreme opinions are unstable, thereby, making this model the opposite of the model with stubborn extremists:

$$\dot{x} = -\text{diag}(x)^2Lx. \quad (3.4)$$

This model assumes that the neutral opinion 0 correspond to a social norm, and the agents may not feel comfortable deviating from it and going against what is acceptable in their society.

Stating our specialized models, we have provided three particular definitions of the agents' susceptibility $A(x)$. While these definitions agree with several well-established socio-psychological theories, the latter theories do not provide any specifics about the particular mathematical form of $A(x)$, besides giving a general idea of its behavior. In our definitions, we use low-degree polynomials, making sure $A(x)$ fits the socio-psychological theories and, at the same time, is simple enough to allow a clear analysis. Similarly, quadratic polynomials were used by Taylor [158] who extended Abelson's linear models [49] with "variable resistance". Alternatively, Friedkin in his recent study [101], when defining the constant susceptibility $A(x(0))$ of the agents based on their initial beliefs, used functions of similar "shape", yet, expressed them using exponential functions.

Nevertheless, despite this lack of a single correct mathematical form for each version of $A(x)$, the convergence results we obtain in the next Section 3.5 are derived independently of a particular mathematical form of $A(x)$, only relying on the zeros of $A(x)$ as well as our ability to analytically compute them. Thus, while we will further provide our analysis for the specialized models using the particular susceptibility functions defined above, this analysis is easy to adapt to other susceptibility functions behaving similarly, yet, possibly, having different mathematical form.

3.4 Preliminaries from Non-smooth Analysis

Prior to analyzing our models for polar opinion dynamics defined in the previous Section 3.3, we will, first, introduce necessary preliminaries. More specifically, we review several tools from non-smooth analysis that prove useful when dealing with non-smooth Lyapunov functions in the proofs of convergence in the analysis of our models in Section 3.5.

Definition 3 (Locally Lipschitz function [159, p.9]). *Function $V : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is locally Lipschitz on $D \subseteq \mathbb{R}^n$, if for each $x \in D$, there exists a constant $K(x) \geq 0 \in \mathbb{R}$, such that for all y in a small-enough ball $\mathfrak{B}(x, \epsilon)$ of radius ϵ centered at x ,*

$$\|V(x) - V(y)\| \leq K(x)\|x - y\|.$$

Definition 4 (Regular function [159, pp.39–40]). *Given $V : \mathbb{R}^n \rightarrow \mathbb{R}$, as well as its right directional derivative at $x \in \mathbb{R}^n$ in the direction of $d \in \mathbb{R}^n$*

$$V'(x; d) = \lim_{h \rightarrow +0} \frac{V(x + hd) - V(x)}{h},$$

and its generalized directional derivative at x in the direction of d

$$V^\circ(x; d) = \lim_{\substack{\delta \rightarrow +0 \\ \epsilon \rightarrow +0}} \sup_{\substack{y \in \mathfrak{B}(x, \delta) \\ h \in [0, \epsilon)}} \frac{V(y + hd) - V(y)}{h},$$

if, for all directions $d \in \mathbb{R}^n$, $V'(x; d)$ exists, and $V'(x; d) = V^\circ(x; d)$, then V is regular at x . If V is convex or/and continuously-differentiable, then it is regular.¹

Definition 5 ((Clarke) Generalized gradient [159]). *Let $V : \mathbb{R}^n \rightarrow \mathbb{R}$ be locally Lipschitz, and Ω_V be the set of points where V fails to be differentiable. Then, the generalized gradient of V is defined as follows*

$$\partial V(x) = \text{co} \left\{ \lim_{i \rightarrow \infty} \nabla V(x_i) \mid x_i \rightarrow x, x_i \notin Z \cup \Omega_V \right\},$$

where $\text{co}\{\cdot\}$ is the convex hull, and Z is any set of Lebesgue measure zero. Thus, the generalized gradient of V at x is the convex hull of all gradient values around and approaching x where V is differentiable.

¹It is a sufficient condition for regularity, not included in the original definition of a regular function in [159].

Theorem 6 (Properties of generalized gradient [160]). *Let $V_1, V_2 : \mathbb{R}^n \rightarrow \mathbb{R}$ be locally Lipschitz and regular at $x \in \mathbb{R}^n$, $a, b \in [0, \infty) \subset \mathbb{R}$. Then,*

- (i) [scaling rule] $\partial(a \cdot V_1)(x) = a \cdot \partial V_1(x)$ and $(a \cdot V_1)$ is locally Lipschitz and regular at x .
- (ii) [sum rule] $\partial(a \cdot V_1 + b \cdot V_2)(x) = a\partial V_1(x) + b\partial V_2(x)$, and $a \cdot V_1 + b \cdot V_2$ is locally Lipschitz and regular at x . The sum of sets in the expression above is understood in the sense of $A + B = \{a + b \mid a \in A, b \in B\}$.
- (iii) [max-min functions] Let $V_k : \mathbb{R}^n \rightarrow \mathbb{R}$, $k \in \{1, \dots, m\} < \infty$ be locally Lipschitz at $x \in \mathbb{R}^n$, and

$$V_{\max}(y) \triangleq \max \{V_k(y) \mid k \in \{1, \dots, m\}\},$$

$$V_{\min}(y) \triangleq \min \{V_k(y) \mid k \in \{1, \dots, m\}\}.$$

Also, let

$$I_{\max}(x) = \{i \mid V_i(x) = V_{\max}(x)\},$$

$$I_{\min}(x) = \{i \mid V_i(x) = V_{\min}(x)\}.$$

Then,

- V_{\max} and V_{\min} are locally Lipschitz at x .
- If V_i is regular at x for each $i \in I_{\max}(x)$, then

$$\partial V_{\max}(x) = \text{co} \cup \{\partial V_i(x) \mid i \in I_{\max}(x)\}$$

and V_{\max} is regular at x .

- If $-V_i$ is regular at x for each $i \in I_{\min}(x)$, then

$$\partial V_{\min}(x) = \text{co} \cup \{\partial V_i(x) \mid i \in I_{\min}(x)\}$$

and $-V_{\min}$ is regular at x .

The following theorem is as a corollary of Theorem 6.

Theorem 7 (Generalized gradients of max-min functions). *Consider functions*

$$V_{\max}(x) = \max(x), \quad V_{-\min}(x) = -\min(x), \quad V_{\max-\min}(x) = V_{\max}(x) + V_{-\min}(x),$$

where $x \in S \subseteq [-1, 1]^n$. Also, let $N = S \cap \{\alpha \mathbf{1} \mid \alpha \in [-1, 1]\}$, and define $I_{\max}(x)$ and $I_{\min}(x)$ as in Theorem 6, and $I_{\text{mid}}(x) = \{1, \dots, n\} \setminus I_{\max}(x) \setminus I_{\min}(x)$. Then,

$$\begin{aligned} \partial V_{\max}(x) &= \{P^\top[\alpha^\top, \mathbf{0}^\top]^\top\}, \\ \partial V_{-\min}(x) &= \{P^\top[-\beta^\top, \mathbf{0}^\top]^\top\}, \\ \partial V_{\max-\min}(x) &= \begin{cases} \{P^\top[\alpha^\top, -\beta^\top, \mathbf{0}^\top]^\top\}, & \text{if } x \in S \setminus N, \\ \{P^\top(\alpha - \beta)\}, & \text{if } x \in N, \end{cases} \end{aligned}$$

where α and β are vectors whose elements comprise coefficients of convex combinations ($\alpha_i, \beta_j \geq 0, \sum_i \alpha_i = \sum_j \beta_j = 1$), α_i correspond to the agents from $I_{\max}(x)$, β_j correspond to the agents from $I_{\min}(x)$, $\mathbf{0}_k$ correspond to the agents from $I_{\text{mid}}(x)$, and permutation matrices P^\top restore the original order of $\{x_i\}$.

Proof. Notice that both $V_{\max}(x) = \max(x)$ and $V_{\min}(x) = \min(x)$ can be viewed as, respectively, the maximum and the minimum of a finite number of functions $V_i(x) = x_i, i \in \{1, \dots, n\}$. Since each $V_i(x)$ is continuously-differentiable and, thus, locally Lipschitz and, due to Defini-

tion 4, regular on S , Theorem 6 allows us to apply the rule (iii) for computing the generalized gradient for max-min functions, followed by the application of the (i) scaling and (ii) sum rules. The statement of the theorem, then, follows immediately. ■

Definition 6 (Set-valued Lie derivative [161, 160]). *For a locally Lipschitz $V : \mathbb{R}^n \rightarrow \mathbb{R}$ and system $\dot{x} = f(x)$, the set-valued Lie derivative $\widetilde{\mathcal{L}}_f V(x)$ of V along the trajectories of the system is defined as $\widetilde{\mathcal{L}}_f V(x) = \{a \in \mathbb{R} \mid \forall \xi \in \partial V(x) : \langle \xi, f(x) \rangle = a\}$.*

The following theorem is an analog of the generalized Invariance Principle [161, 160], specialized for the case of a continuous vector field, while the original was stated for differential inclusions.

Theorem 8 (Invariance Principle [161, 160]). *If*

(i) $V : \mathbb{R}^n \rightarrow \mathbb{R}$ *is locally Lipschitz and regular,*

(ii) $S \subset \mathbb{R}^n$ *is compact and forward invariant with respect to $\dot{x} = f(x)$, and*

(iii) $\max \widetilde{\mathcal{L}}_f V(x) \leq 0$ *for each $x \in S$, then all solutions $x(t) : [0, \infty) \rightarrow \mathbb{R}^n$ starting in S converge to the largest invariant subset M of $S \cap \overline{\{x \in \mathbb{R}^n \mid 0 \in \widetilde{\mathcal{L}}_f V(x)\}}$, where $\overline{\{\dots\}}$ is set closure. If M is finite, then the limit of each solution $x(0) \in S$ exists and is an element of M .*

3.5 Analysis of the Models

In this section, we will analyze well-posedness, equilibrium points, and the asymptotic behavior of our models. The convergence proofs will rely on several notions from non-smooth analysis reviewed in the previous Section 3.4.

3.5.1 Well-posedness

In order for our general model $\dot{x} = -A(x)Lx$ to be well-posed, its solutions must exist, be unique, and must never escape the state space $[-1, 1]^n$. These properties of the solutions are stated in the following theorem and its corollary.

Theorem 9 (Well-posedness of the general model). *If $x(0) \in [-1, 1]^n$, and the evolution of $x(t)$ is governed by the general model of polar opinion dynamics $\dot{x} = -A(x)Lx$, then $x(t) \in [-1, 1]^n$ for any $t \geq 0$.*

Since Theorem 9 is a standard result in control theory, we omit its proof and only mention that its validity immediately follows from the general contraction Lemma 5.

Corollary (Existence, uniqueness, smoothness of solutions). *The general model of polar opinion dynamics $\dot{x} = -A(x)Lx$, $x(0) \in [-1, 1]^n$, has a unique continuously-differentiable solution $x(t)$ defined for all $t > 0$.*

Proof. We will use Theorem 3.3 of [162], according to which, if $f(x)$ is locally Lipschitz in x in the entire compact $[-1, 1]^n$, and each solution $x(t)$ of (3.1) lies entirely in $[-1, 1]^n$, then (3.1) has a unique solution defined for all $t > 0$. The latter property of the solutions comes from the proven above Theorem 9, and what remains is to show that $f(x)$ is locally Lipschitz in x at each point of $[-1, 1]^n$.

From the definition (3.1) of the general model, $A(x)$ is locally Lipschitz in x at each point of $[-1, 1]^n$, that is, for all $x \in [-1, 1]^n$, there is $K(x) \geq 0 \in \mathbb{R}$, such that for all y in a small enough ball $\mathfrak{B}(x, \epsilon)$ of radius ϵ centered at x , we have

$$\|A(x) - A(y)\| \leq K(x)\|x - y\|. \quad (3.5)$$

In order to obtain a similar inequality for $f(x)$, we proceed as follows, having defined matrix norm $\|\cdot\|$ to be an operator norm corresponding to ℓ_2 vector norm.

$$\begin{aligned}
\|f(x) - f(y)\| &= \|A(x)Lx - A(y)Ly\| \\
&= \|A(x)Lx - A(y)Lx + A(y)Lx - A(y)Ly\| \\
&= \|(A(x) - A(y))Lx + A(y)L(x - y)\| \\
&\leq \|A(x) - A(y)\| \cdot \|Lx\| + \|A(y)\| \cdot \|L\| \cdot \|x - y\| \\
&\leq K(x)\|x - y\| \cdot \|L\| \cdot \|x\| + \|A(y)\| \cdot \|L\| \cdot \|x - y\| \\
&= \|L\|(K(x)\|x\| + \|A(y)\|) \cdot \|x - y\| = R(x)\|x - y\|.
\end{aligned}$$

In the obtained expression for $R(x)$, $\|L\| < \infty$ is a norm of a finite matrix, $K(x)$ is a Lipschitz constant of $A(x)$, $\|x\| = \|x\|_2 \leq \sqrt{n} < \infty$ in $[-1, 1]^n$, and $\|A(y)\| < \infty$ as $A(x)$ is locally Lipschitz and, hence, bounded in a compact $[-1, 1]^n$. Thus, $R(x)$ is a Lipschitz constant for $f(x)$, and, hence, $f(x)$ is locally Lipschitz in x at each $x \in [-1, 1]^n$.

As to the smoothness of $x(t)$, since $A(x)$ is locally Lipschitz, then $f(x)$ is continuous, and, hence, $x(t)$ is continuously-differentiable. ■

3.5.2 Equilibrium Points

Prior to studying the equilibrium points of our models, we will prove a basic lemma.

Lemma 3 (Properties of some network partitions). *Let $W \in \mathbb{R}^{n \times n}$ be a row-stochastic adjacency matrix of a strongly connected network $G(W)$. If $G(W)$'s nodes are partitioned into two non-empty sets $\{1, \dots, n\} = I_1 \cup I_2$, $I_1 \cap I_2 = \emptyset$, and P is any permutation matrix such that*

nodes I_1 precede nodes I_2 in

$$PWP^\top = \begin{array}{c} I_1 \quad I_2 \\ \begin{bmatrix} W_{11} & W_{12} \\ W_{21} & W_{22} \end{bmatrix} \end{array},$$

then both $(I - W_{11})$ and $(I - W_{22})$ are invertible, and both $(I - W_{11})^{-1}W_{12}\mathbb{1} = \mathbb{1}$ and $(I - W_{22})^{-1}W_{21}\mathbb{1} = \mathbb{1}$.

Proof. Since $G(W)$ is strongly connected, W is irreducible. Consequently, since both I_1 and I_2 are non-empty, W_{11} is substochastic, so there exists $\ell \in I_1$, such that

$$\sum_{j \in I_1} (W_{11})_{\ell j} < 1. \quad (*)$$

Notice that $\sum_{j \in I_1} (W_{11})_{ij}^k$ can be interpreted as the likelihood of a k -hop random walk on $G(W)$ to start at node i and end at any node of I_1 , so $(*)$ implies that there is a positive likelihood for a 1-hop random walk starting at ℓ to escape I_1 . If we define $d(i, j)$ to be the length—in hops—of the shortest path from node i to node j in $G(W)$, and $d_{\max}(\ell) = \max_{i \in I_1} d(i, \ell)$, then

$$\forall k > d_{\max}(\ell) \forall i \in I_1 : \sum_{j \in I_1} (W_{11})_{ij}^k < 1,$$

since for each $i \in I_1$, there is at least one k -hop walk passing through ℓ , and, as a result, there is a positive likelihood of any such walk's escaping I_1 . Hence, for all $k > d_{\max}(\ell)$, W_{11}^k is convergent, and its spectral radius $\rho(W_{11}^k) < 1$, which immediately entails $\rho(W_{11}) < 1$. Hence, for the spectrum of $(I - W_{11})$, we have $\sigma(I - W_{11}) \subset (0, 2)$. Thus, $(I - W_{11})$ is non-singular and,

as such, invertible. Consequently, matrix $(I - W_{11})^{-1}W_{12}$ is well-defined, and its row-sums are

$$\begin{aligned} (I - W_{11})^{-1}W_{12}\mathbf{1} &= (\text{since } W \text{ is row-stochastic}) \\ &= (I - W_{11})^{-1}(\mathbf{1} - W_{11}\mathbf{1}) = (I - W_{11})^{-1}(I - W_{11})\mathbf{1} = \mathbf{1}. \end{aligned}$$

Applying the same reasoning to blocks W_{22} and W_{21} in place of blocks W_{11} and W_{12} , we obtain the existence of $(I - W_{22})^{-1}$, and equality $(I - W_{22})^{-1}W_{21}\mathbf{1} = \mathbf{1}$. ■

Theorem 10 (Equilibrium points). *Suppose the network's adjacency matrix W is row-stochastic, and network $G(W)$ is strongly connected. Then, the following holds.*

1) *The equilibrium points of the stubborn positives model $\dot{x} = -\frac{1}{2}(I - \text{diag}(x))Lx$ and the stubborn neutrals model $\dot{x} = -\text{diag}(x)^2Lx$ are*

$$x^* = \alpha\mathbf{1}_n, \quad \alpha \in [-1, 1].$$

2) *Consider an arbitrary agent set partition $\{1, \dots, n\} = I_1 \cup I_2$, $I_1 \cap I_2 = \emptyset$, $2 \leq |I_1| \leq n$, and an arbitrary permutation matrix P such that the agents are ordered as*

$$PW P^\top = \begin{array}{c} I_1 \quad I_2 \\ I_1 \begin{bmatrix} W_{11} & W_{12} \\ W_{21} & W_{22} \end{bmatrix} \end{array}, \quad Px = \begin{array}{c} I_1 \\ I_2 \end{array} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}.$$

Then, the equilibrium points of the stubborn extremists model $\dot{x} = -(I - \text{diag}(x)^2)Lx$ are

$$x^* = \alpha\mathbf{1}_n, \quad \alpha \in [-1, 1], \quad \text{and} \quad x^* = P^\top [x_1^{*\top}, x_2^{*\top}]^\top,$$

where

$$x_1^* \in \{-1, 1\}^{|I_1|} \setminus \{-\mathbb{1}_{|I_1|}, \mathbb{1}_{|I_1|}\},$$

$$x_2^* = \begin{cases} (I - W_{22})^{-1} W_{21} x_1^*, & \text{if } I_2 \neq \emptyset, \\ [\]_{0 \times 1}, & \text{if } I_2 = \emptyset. \end{cases}$$

Proof. 1) First, let us deal with the model with stubborn positives

$$\dot{x} = f(x) = -\frac{1}{2}(I - \text{diag}(x))Lx.$$

It is easy to see that $x^* = \alpha \mathbb{1}$, $\alpha \in [-1, 1]$ are equilibrium points of the system. Now, let us look for equilibrium points corresponding to the states of disagreement. Consider such a candidate point $x \in [-1, 1]^n$, $x \neq \alpha \mathbb{1}$, $n > 1$. Since $x \neq \alpha \mathbb{1}$, there exists agent $i \in \{1, \dots, n\}$, such that $x_i = \min(x) < 1$, and exists agent $j \in \{1, \dots, n\}$ such that $x_j > x_i$ and $W_{ij} > 0$. Because $x_i < 1$, $A_{ii}(x) = \frac{1}{2}(1 - x_i) > 0$; due to the existence of such agent j , $(Lx)_i \neq 0$. As a result, $(f(x))_i \neq 0$ and, thus, x is not a point of equilibrium. Hence, $x^* = \alpha \mathbb{1}$ are the only points of equilibrium of the model with stubborn positives.

In the remainder of the proof, we will consider different partitions $\{1, \dots, n\} = I_1 \cup I_2$, $I_1 \cap I_2 = \emptyset$ of the agent set, and P being any permutation matrix such that agents I_1 precede agents I_2 in $PW P^\top$ and Px . For readability, for each partition $I_1 \cup I_2$, we will omit P in the expressions for W and x , and will apply the right agent ordering later.

Let us proceed to the model with stubborn neutrals

$$\dot{x} = -\text{diag}(x)^2 Lx.$$

A candidate equilibrium point x is defined with respect to an agent set partition $\{1, \dots, n\} =$

$I_1 \cup I_2, I_1 \cap I_2 = \emptyset, 0 \leq |I_1| \leq n$ as follows: $x_i = 0$ if $i \in I_1$; and $|x_i| > 0$ if $i \in I_2$. If $I_1 = \emptyset$, then $f(x) = 0 \Leftrightarrow Lx = 0 \Leftrightarrow x^* = \alpha \mathbb{1}$, where $\alpha \in [-1, 1]$. If $I_2 = \emptyset$, then, clearly, $x^* = 0$. Finally, if both I_1 and I_2 are non-empty, then, w.l.o.g., assuming that agents I_1 precede agents I_2 , $f(x) = 0 \Leftrightarrow (I - W_{22})x_2 = 0$. However, from Lemma 3, we know that $(I - W_{22})$ is invertible. Hence, the obtained equation has only the trivial solution $x_2 = 0$, and the corresponding equilibrium point is $x^* = 0$. We have proven that the equilibrium points of the model with stubborn neutrals are $x^* = \alpha \mathbb{1}, \alpha \in [-1, 1]$.

2) For the model with stubborn extremists

$$\dot{x} = -(I - \text{diag}(x)^2)Lx,$$

we define candidate equilibrium points with respect to partition $\{1, \dots, n\} = I_1 \cup I_2, I_1 \cap I_2 = \emptyset, 0 \leq |I_1| \leq n$ as follows: $|x_i| = 1$ if $i \in I_1$; and $|x_i| < 1$ if $i \in I_2$. If $I_1 = \emptyset$, then $f(x) = 0 \Leftrightarrow Lx = 0$, and, thus, $x^* = \alpha \mathbb{1}, \alpha \in [-1, 1]$. If $I_2 = \emptyset$, then $x^* \in \{-1, 1\}^n$. If both I_1 and I_2 are non-empty, then, again, assuming that agents I_1 precede agents I_2 in W and x , equation $f(x) = 0$ is rewritten as

$$(I - \text{diag}(x_2)^2) \begin{bmatrix} -W_{21} & (I - W_{22}) \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = 0$$

$$\Leftrightarrow (\text{since } |x_2| < 1) \Leftrightarrow (I - W_{22})x_2 = W_{21}x_1.$$

From Lemma 3, we know that $(I - W_{22})$ is invertible. Thus, $x_2 = (I - W_{22})^{-1}W_{21}x_1$. Among the $x = [x_1^\top, x_2^\top]^\top$ satisfying the obtained equation, we would like to separate those corresponding to consensus and those corresponding to disagreement. If $x_1 = \mathbb{1}$, then, again from Lemma 3, we derive $x_2 = \mathbb{1}$, and $x^* = \mathbb{1}$. Similarly, if $x_1 = -\mathbb{1}$, then $x^* = -\mathbb{1}$. Notice, that all the equilibrium points discovered so far correspond to consensus and are independent of matrix P .

Finally, if $x_1 \in \{-1, 1\}^{|I_1|} \setminus \{\mathbb{1}, -\mathbb{1}\}$ (which implies $|I_1| \geq 2$), and, hence, the agents necessarily disagree, then $x_2 = (I - W_{22})^{-1}W_{21}x_1$, and the corresponding equilibrium points, under the partition-defined agent order P , are $x^* = P^\top \begin{bmatrix} x_1^\top & x_2^\top \end{bmatrix}^\top$. ■

3.5.3 Convergence Analysis

Having studied the equilibrium points of our specialized models, we will now study these models' convergence. We will, first, establish sufficient conditions for convergence to consensus of the general model of polar opinion dynamics and, then, use this result to prove convergence of the three specialized models.

In the proofs of convergence, we will need establishing forward invariance of certain subsets of the state space with respect to a model at hand. To that end, we will need the following lemma, being an immediate consequence of the solution uniqueness stated in the Corollary of Theorem 9.

Lemma 4 (Agent subset invariance). *If $x(t) \in [-1, 1]^n$ evolves according to one of the specialized models*

$$\dot{x} = -\frac{1}{2}(I - \text{diag}(x))Lx, \quad (3.3)$$

$$\dot{x} = -\text{diag}(x)^2Lx, \quad (3.4)$$

$$\dot{x} = -(I - \text{diag}(x)^2)Lx, \quad (3.2)$$

and the agents are partitioned into $I_{\text{closed}}(t) = \{i \mid A_{ii}(x(t)) = 0\}$ and $I_{\text{open}}(t) = \{i \mid A_{ii}(x(t)) > 0\}$, then, for all $t \geq 0$, $I_{\text{closed}}(t) = I_{\text{closed}}(0) = I_{\text{closed}}$ and $I_{\text{open}}(t) = I_{\text{open}}(0) = I_{\text{open}}$.

The following lemma will be instrumental in proving forward invariance of subsets of the state space as well as in the construction of Lyapunov functions in the convergence proofs.

Lemma 5 (General contraction lemma). *Suppose that W is a row-stochastic adjacency matrix*

of the network, and agent states $x(t) \in [-1, 1]^n$ evolve according to the general model of polar opinion dynamics

$$\dot{x} = f(x) = -A(x)Lx. \quad (3.1)$$

Then, $V_{\max}(x) = \max(x)$ is non-increasing and $V_{\min}(x) = \min(x)$ is non-decreasing along the trajectories of (3.1).

Proof. Let us consider $V_{\max}(x) = \max(x)$ and define $I_{\max}(x) = \{i \mid x_i = \max(x)\}$. According to Lemma 2.2 of Lin et al. [163], the upper Dini derivative of V_{\max} along the trajectories of (3.1) is defined as

$$D_f^+ V_{\max}(x) = \max_{i \in I_{\max}(x(t))} \dot{x}_i(t) = \max_{i \in I_{\max}(x(t))} \underbrace{A_{ii}(x(t))}_{\in [0,1]} \sum_{j \in N^{out}(i)} \underbrace{w_{ij}(x_j - x_i)}_{\leq 0} \leq 0.$$

Hence, V_{\max} is non-increasing along the trajectories of (3.1). The proof for V_{\min} is similar and, hence, is omitted. ■

We have laid out all the necessary preliminaries, and are ready to prove convergence of our models. In the following theorem, we will establish a sufficient condition for the convergence to consensus of the general model of polar opinion dynamics.

Theorem 11 (General convergence to consensus). *Suppose that W is a row-stochastic adjacency matrix of a strongly connected network $G(W)$, and agent states $x(t) \in [-1, 1]^n$ evolve according to the general model of polar opinion dynamics*

$$\dot{x} = f(x) = -A(x)Lx, \quad (3.1)$$

with the agents' having potentially different susceptibility functions $A_{ii}(x)$. Let $S \subseteq [-1, 1]^n$ be

a non-empty compact set, forward invariant with respect to system (3.1), and

$$N = S \cap \{\alpha \mathbf{1} \mid \alpha \in [-1, 1]\}$$

be its non-empty subset of consensus states. Further, assume that in S , the agents' susceptibility functions $A_{ii}(x)$ agree upon their zeros in that

$$\forall x \in S \forall i, j \in \{1, \dots, n\} : A_{ii}(x) = A_{jj}(x) = 0 \rightarrow x_i = x_j.$$

Then, all trajectories $x(t)$ of (3.1) starting in S converge to N as $t \rightarrow \infty$.

Proof. We will prove convergence using the Invariance Principle given as Theorem 8 in Section 3.4. To apply it, we will, first, need to find a suitable Lyapunov function for system (3.1). Consider the following Lyapunov function candidate

$$V_{\max\text{-min}}(x) = \max(x) - \min(x).$$

While it immediately follows from the general contraction Lemma 5 that $V_{\max\text{-min}}(x)$ is non-increasing along the trajectories of system (3.1), in order to prove convergence to a set, we need a more detailed analysis that would distinguish the cases when $V_{\max\text{-min}}(x)$ decreases and when it does not change along the system's trajectories.

From Theorem 7, it follows that

$$\partial V_{\max\text{-min}}(x) = \begin{cases} \{P^\top[\alpha^\top, -\beta^\top, \mathbf{0}^\top]^\top\}, & \text{if } x \in S \setminus N, \\ \{P^\top(\alpha - \beta)\}, & \text{if } x \in N, \end{cases}$$

where convex combination coefficients α_i correspond to the agents in $I_{\max}(x) = \{i \mid x_i = \max(x)\}$, convex combination coefficients β_i correspond to the agents in $I_{\min}(x) = \{i \mid x_i =$

$\min(x)\}$, \emptyset correspond to the rest of the agents $I_{\text{mid}}(x) = \{1, \dots, n\} \setminus I_{\text{max}}(x) \setminus I_{\text{min}}(x)$, and permutation matrix P^T restores the original agent order.

Now, for each $\xi \in \partial V_{\text{max-min}}(x)$, we are interested in the values of inner products $\langle \xi, f(x) \rangle$, which, according to Definition 6, comprise the set-valued Lie derivative

$$\widetilde{\mathcal{L}}_f V_{\text{max-min}}(x) = \{a \in \mathbb{R} \mid \forall \xi \in \partial V_{\text{max-min}}(x) : \langle \xi, f(x) \rangle = a\}$$

of $V_{\text{max-min}}(x)$ at x along the trajectories of system (3.1). Our immediate goal is to understand when $\max \widetilde{\mathcal{L}}_f V_{\text{max-min}}(x)$ is negative and when it is zero, depending on the chosen $x \in S$.

If $x \in N$, that is, $x = \alpha \mathbb{1}$ for some $\alpha \in [-1, 1]$, then $f(x) = -A(x)Lx = -A(x)\alpha(L\mathbb{1}) = \emptyset$ and, thus, $\forall \xi \in \partial V_{\text{max-min}}(x) : \langle \xi, f(x) \rangle = 0$, so $\widetilde{\mathcal{L}}_f V_{\text{max-min}}(x) = \{0\}$.

If $x \in S \setminus N$, then let us investigate the possible values of $\langle \xi, f(x) \rangle$, w.l.o.g., dropping P in the expression for ξ , for readability, and using the same agent order in $f(x)$ as in ξ :

$$\begin{aligned} \langle \xi, f(x) \rangle &= -\xi^T A(x) L x \\ &= - \begin{bmatrix} \alpha \\ -\beta \\ 0 \end{bmatrix}^T \begin{bmatrix} A_{\text{max}}(x) & 0 & 0 \\ 0 & A_{\text{min}}(x) & 0 \\ 0 & 0 & A_{\text{mid}}(x) \end{bmatrix} \begin{bmatrix} (I - W_{11}) & -W_{12} & -W_{13} \\ -W_{21} & (I - W_{22}) & -W_{23} \\ -W_{31} & -W_{32} & (I - W_{33}) \end{bmatrix} \begin{bmatrix} x_{\text{max}} \\ x_{\text{min}} \\ x_{\text{mid}} \end{bmatrix} \\ &= -(\alpha^T A_{\text{max}}(x)(x_{\text{max}} - [W_{11}W_{12}W_{13}]x) + \beta^T A_{\text{min}}(x)([W_{21}W_{22}W_{23}]x - x_{\text{min}})), \end{aligned}$$

where x_{max} , x_{min} , and x_{mid} are the states of the agents from $I_{\text{max}}(x)$, $I_{\text{min}}(x)$, and $I_{\text{mid}}(x)$, respectively; $A_{\text{max}}(x)$, $A_{\text{min}}(x)$, and $A_{\text{mid}}(x)$ are the diagonal matrices of susceptibilities of the agents from these three agent subsets; and adjacency matrix W is partitioned ordering the agents as $I_{\text{max}}(x)$, $I_{\text{min}}(x)$, $I_{\text{mid}}(x)$.

Since $x \in S \setminus N$, then $x \neq \alpha \mathbb{1}$ and $\max(x) > \min(x)$. Hence, from the assumption of the theorem about the agreement of $A_{ii}(x)$ upon zeros, it follows that at least one of the inequalities

$\text{diag}(A_{\max}(x)) > \mathbb{0}$ and $\text{diag}(A_{\min}(x)) > \mathbb{0}$ holds. Let us assume, for now, that $\text{diag}(A_{\max}(x)) > \mathbb{0}$ and focus on the first term $\alpha^\top A_{\max}(x)(x_{\max} - [W_{11}W_{12}W_{13}]x)$ of the obtained expression for $\langle \xi, f(x) \rangle$.

Due to row-stochasticity of W and strong connectivity of $G(W)$, for equality $(x_{\max} - [W_{11}W_{12}W_{13}]x)_i = 0$ to hold, all out-neighbors of agent $i \in I_{\max}(x)$ in $G(W)$ must also be from $I_{\max}(x)$. If there are no such agents i , then $x_{\max} - [W_{11}W_{12}W_{13}]x > \mathbb{0}$. Additionally, since α is comprised of a convex combination's coefficients, and $\text{diag}(A_{\max}(x)) > \mathbb{0}$, then at least one element of $\alpha^\top A_{\max}(x)$ is positive. Hence, $\alpha^\top A_{\max}(x)(x_{\max} - [W_{11}W_{12}W_{13}]x) > 0$, and $\max \widetilde{\mathcal{L}}_f V_{\max\text{-min}}(x) < 0$.

If there is an agent $i \in I_{\max}(x)$ with its entire out-neighborhood consisting of the members of $I_{\max}(x)$, then $(x_{\max} - [W_{11}W_{12}W_{13}]x)_i = 0$. However, $x_{\max} - [W_{11}W_{12}W_{13}]x \neq \mathbb{0}$, as the opposite would be possible either if agents $I_{\max}(x)$ were disconnected from the rest of the network (which is impossible due to the network's strong connectivity assumption), or x was a consensus state (which is impossible, as such states are absent from $S \setminus N$). Thus, there is $j \in I_{\max}(x)$ such that $(x_{\max} - [W_{11}W_{12}W_{13}]x)_j = \delta > 0$. Now, however, if we put $\alpha_1 = e_i$ and $\alpha_2 = (e_i + e_j)/2$, with e_k being the k 'th element of the standard basis, we will have

$$\alpha_1^\top A_{\max}(x)(x_{\max} - [W_{11}W_{12}W_{13}]x) = 0,$$

$$\alpha_2^\top A_{\max}(x)(x_{\max} - [W_{11}W_{12}W_{13}]x) = \delta/2(A_{\max}(x))_{jj} > 0.$$

Consequently, for a given $x \in S \setminus N$, term $\alpha^\top A_{\max}(x)(x_{\max} - [W_{11}W_{12}W_{13}]x)$ takes at least two different values, depending on the choice of α . It can be analogously shown that, if $\text{diag}(A_{\min}(x)) > \mathbb{0}$, then $\beta^\top A_{\min}(x)([W_{21}W_{22}W_{23}]x - x_{\min})$ also takes at least two different values, for different β . Hence, since α and β can be chosen independently, and at least one of the terms $\alpha^\top A_{\max}(x)$ and $\beta^\top A_{\min}(x)$ is not $\mathbb{0}$, we conclude that, if $x \in S \setminus N$, and there are some

agents in $I_{\max}(x)$ whose entire out-neighborhood is also in $I_{\max}(x)$, then

$$\exists \xi_1 \neq \xi_2 : \langle \xi_1, f(x) \rangle \neq \langle \xi_2, f(x) \rangle,$$

which entails $\widetilde{\mathcal{L}}_f V_{\max\text{-min}}(x) = \emptyset$ and, by convention, $\max \widetilde{\mathcal{L}}_f V_{\max\text{-min}}(x) = \max \emptyset = -\infty < 0$.

To summarize, we have so far shown that, if $x \in S \setminus N$, then $\max \widetilde{\mathcal{L}}_f V_{\max\text{-min}}(x) < 0$, and, if $x \in N$, then $\widetilde{\mathcal{L}}_f V_{\max\text{-min}}(x) = \{0\}$. Additionally, it immediately follows from Theorem 6 that $V_{\max\text{-min}}(x)$ is Lipschitz and regular on S . Thus, $V_{\max\text{-min}}(x)$ is a Lyapunov function for system (3.1).

Finally, we notice that, by assumption, S is compact and forward invariant with respect to system (3.1). Additionally, N , in which $0 \in \widetilde{\mathcal{L}}_f V_{\max\text{-min}}(x)$, is forward invariant with respect to system (3.1)—as it entirely consists of equilibrium points—and, clearly, is the largest closed subset of itself. These two facts, taken together with the existence of Lyapunov function $V_{\max\text{-min}}(x)$, allow us to conclude that, by Invariance Principle, all trajectories $x(t)$ of system (3.1) starting in S converge to N as $t \rightarrow \infty$. ■

Having proven a sufficient condition for the convergence to consensus of the general model, we will proceed with a comprehensive analysis of convergence for the three specialized models, starting with the model with stubborn positives.

Theorem 12 (Convergence—Stubborn Positives). *Suppose that W is a row-stochastic adjacency matrix of a strongly connected network, and $x(t)$ evolves according to the model with stubborn positives*

$$\dot{x} = f(x) = -A(x)Lx, \quad A(x) = \frac{1}{2}(I - \text{diag}(x)). \quad (3.3)$$

Then,

- if $x(0) < \mathbb{1}$, then $\lim_{t \rightarrow \infty} x(t) = \alpha \mathbb{1}$, $\alpha \in [-1, 1)$;

- if exists i such that $x_i(0) = 1$, then $\lim_{t \rightarrow \infty} x(t) = \mathbb{1}$.

In other words, in the absence of the agents initially having extreme states, $x(t)$ converges to some consensus state $\alpha \mathbb{1}$, $\alpha < 1$, and if there is at least one agent initially holding the extreme state of 1, then all agents approach that state as $t \rightarrow \infty$.

Proof. Since the convergence behavior of model (3.3) varies across the state space, let us, first, partition the latter and, then, prove convergence in each part individually. Consider the following state space partition, illustrated with Figure 3.1:

$$[-1, 1]^n = \lim_{\epsilon \rightarrow +0} S_0(\epsilon) \cup S_1,$$

$$S_0(\epsilon) = [-1, 1 - \epsilon]^n,$$

$$S_1 = \{P^T x \mid x \in \cup_{k=1}^n \{1\}^k \times [-1, 1)^{n-k}\},$$

where $N_0(\epsilon) = \{\alpha \mathbb{1} \mid \alpha \in [-1, 1 - \epsilon]\}$ and $N_1 = \{\mathbb{1}\}$ are the sets of consensus states in $S_0(\epsilon)$ and S_1 , respectively, convergence to which is expected.

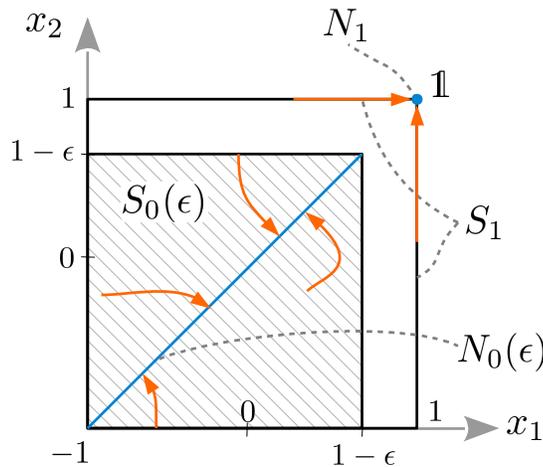


Figure 3.1: Convergence behavior of the model with stubborn positives in two dimensions, as well as the partition of the state space. Several trajectories representative of the model's behavior are displayed as solid arrows.

(i) *Convergence from $S_0(\epsilon)$ to $N_0(\epsilon)$* : We will prove convergence using the general convergence Theorem 11, with $S = S_0(\epsilon)$ and $N = N_0(\epsilon)$. To apply the theorem, we need to prove the agreement upon zeros of the susceptibility functions $A_{ii}(x)$ and forward invariance of $S_0(\epsilon)$. (Theorem 11 also requires both S and N to be non-empty, and S to be compact. However, whenever we use Theorem 11, non-emptiness trivially follows from the definition of these sets, and compactness of S immediately follows from Heine-Borel theorem, as we always choose $S \subseteq \mathbb{R}^n$, $n < \infty$ to be both bounded and closed. Thus, we will further omit the discussion of these two statements about S and N from our proofs.)

Firstly, as $A(x) = \frac{1}{2}(I - \text{diag}(x))$ and, thus, $A_{ii}(x) = \frac{1}{2}(1 - x_i)$, it is clear that, if $A_{ii}(x) = A_{jj}(x) = 0$, then $x_i = x_j = 1$, which proves the zero-agreement property

$$\forall i, j \in 1, \dots, n : A_{ii}(x) = A_{jj}(x) = 0 \rightarrow x_i = x_j.$$

In order to prove forward invariance of $S_0(\epsilon)$ with respect to system (3.3), we notice that, according to contraction Lemma 5, $V_{\max}(x) = \max(x)$ is non-increasing along the trajectories of system (3.3), and, at the same time, from the well-posedness Theorem 9, we know that $\min(x) \geq -1$ for all $x \in [-1, 1]^n$. Consequently, all the trajectories of the system starting inside cube $S_0(\epsilon)$ remain in it as $t \rightarrow \infty$.

Now, by invoking the general convergence Theorem 11, we conclude that all trajectories of system (3.3) starting in $S_0(\epsilon)$ converge to $N_0(\epsilon)$ as $t \rightarrow \infty$.

(ii) *Convergence from S_1 to N_1* : The agreement upon zeros property of $A_{ii}(x)$ has already been proven above. As to forward invariance of S_1 with respect to system (3.3), it follows immediately from Lemma 4 about the invariance of the closed agent subset. Thus, by the general convergence Theorem 11, all trajectories of system (3.3) starting in S_1 converge to N_1 as $t \rightarrow \infty$. ■

Theorem 13 (Convergence—Stubborn Neutrals). *Suppose that W is a row-stochastic adjacency matrix of a strongly connected network, and $x(t)$ evolves according to the model with stubborn neutrals*

$$\dot{x} = f(x) = -A(x)Lx, \quad A(x) = \text{diag}(x)^2. \quad (3.4)$$

Then,

- if $x(0) > 0$, then $\lim_{t \rightarrow \infty} x(t) = \alpha \mathbf{1}$, $\alpha \in (0, 1]$;
- if $x(0) < 0$, then $\lim_{t \rightarrow \infty} x(t) = \alpha \mathbf{1}$, $\alpha \in [-1, 0)$;
- otherwise, $\lim_{t \rightarrow \infty} x(t) = 0$.

In other words, if the initial states of all agents are positive, then $x(t)$ converges to an element-wise positive consensus state; if the initial states are all negative, then the convergence is to a negative consensus; finally, if either there are some closed agents, or some open agents' states have opposite signs, then $x(t)$ converges to 0 as $t \rightarrow \infty$.

Proof. Let us, first, partition the state space—as shown in Figure 3.2—and, then, prove convergence for each part individually.

$$\begin{aligned} [-1, 1]^n &= S_0 \cup \lim_{\epsilon \rightarrow +0} S_-(\epsilon) \cup \lim_{\epsilon \rightarrow +0} S_+(\epsilon), \\ S_-(\epsilon) &= [-1, -\epsilon]^n, \quad N_-(\epsilon) = \{\alpha \mathbf{1} \mid \alpha \in [-1, -\epsilon]\}, \\ S_+(\epsilon) &= [\epsilon, 1]^n, \quad N_+(\epsilon) = \{\alpha \mathbf{1} \mid \alpha \in [\epsilon, 1]\}, \\ S_0 &= \{x \in [-1, 1]^n \mid \prod_{i=1}^n x_i = 0 \vee \exists i, j : \text{sgn}(x_i x_j) = -1\}, \\ N_0 &= \{0\}, \end{aligned}$$

where we expect convergence from $S_-(\epsilon)$, $S_+(\epsilon)$, and S_0 to $N_-(\epsilon)$, $N_+(\epsilon)$, and N_0 , respec-

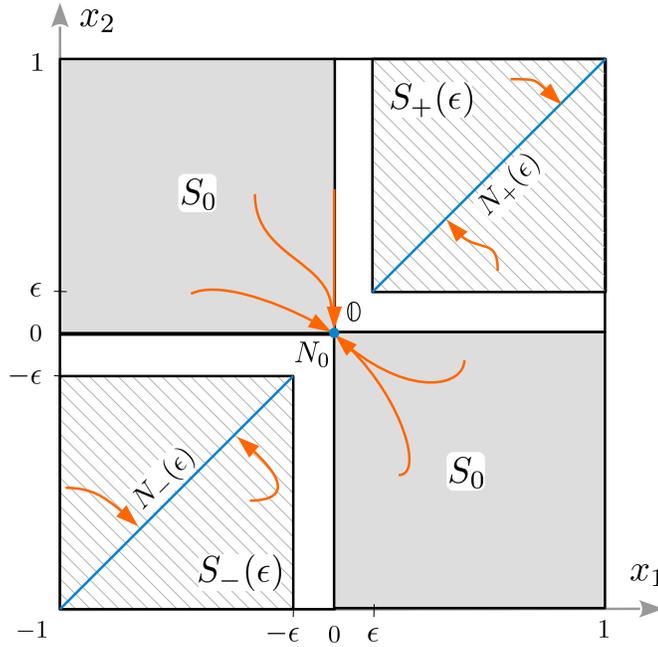


Figure 3.2: Convergence behavior of the model with stubborn neutrals in two dimensions, as well as the partition of the state space. Several trajectories representative of the model's behavior are shown as solid arrows.

tively.

(i) *Convergence from $S_-(\epsilon)$ to $N_-(\epsilon)$* : We will prove convergence using the general convergence Theorem 11, which requires that we prove the agreement upon zeros of functions $A_{ii}(x)$ and forward invariance of $S_-(\epsilon)$.

Since $A(x) = \text{diag}(x)^2$ and $A_{ii}(x) = x_i^2$, it is clear that $A_{ii}(x) = 0 \Leftrightarrow x_i = 0$, thus, proving the zero-agreement property

$$\forall i, j \in 1, \dots, n : A_{ii}(x) = A_{jj}(x) = 0 \rightarrow x_i = x_j.$$

Forward invariance of $S_-(\epsilon)$ immediately follows from the facts that $V_{\max}(x)$ is non-increasing along the trajectories of the system due to the general contraction Lemma 5, and $\max(x) \geq -1$ follows from the well-posedness Theorem 9.

We, now, can invoke the general convergence Theorem 11 and conclude that all trajectories

of (3.4) starting in $S_-(\epsilon)$ converge to $N_-(\epsilon)$ as $t \rightarrow \infty$.

(ii) *Convergence from $S_+(\epsilon)$ to $N_+(\epsilon)$* : The proof is identical to the proof for the case of $S_-(\epsilon)$ and $N_-(\epsilon)$ and, as such, is omitted.

(iii) *Convergence from S_0 to N_0* : The agreement of $A_{ii}(x)$ upon zeros has already been proven in part (i). As to forward invariance of S_0 , there are two qualitatively different ways a trajectory of the system can leave one of the mixed-sign orthants that S_0 consists of: either a trajectory leaves cube $[-1, 1]^n$ or it escapes into either the positive or the negative orthant. The former is impossible due to the well-posedness Theorem 9, which states that the trajectories cannot leave the state space $[-1, 1]^n$. The latter is also impossible, because, in order for a continuous trajectory $x(t)$ to leave from a mixed-sign to the negative or the positive orthant, the closed agent subset $I_{closed}(x)$ has to change when a trajectory passes a mixed-sign orthant's boundary, which would contradict the agent subset invariance Lemma 4. Thus, by the general convergence Theorem 11, all trajectories of (3.4) starting in S_0 converge to N_0 as $t \rightarrow \infty$. ■

Theorem 14 (Convergence—Stubborn Extremists). *Suppose that W is a row-stochastic matrix of a strongly-connected network, and state $x(t)$ is governed by the model with stubborn extremists*

$$\dot{x} = -A(x)Lx, \quad A(x) = (I - \text{diag}(x))^2. \quad (3.2)$$

Further, assume that the agent set is partitioned as

$$I_{open} = \{i \mid A_{ii}(x(0)) > 0\}, \quad I_{closed} = \{i \mid A_{ii}(x(0)) = 0\}.$$

Then, the following holds:

- *If $I_{closed} = \emptyset$, then $\lim_{t \rightarrow \infty} x(t) = \alpha \mathbb{1}$, for some $\alpha \in (-1, 1)$, that is, if there are no closed agents, then the system converges to a consensus.*

- If $I_{closed} \neq \emptyset$, yet, $\forall i, j \in I_{closed} : x_i = x_j = \alpha \in \{-1, 1\}$, then $\lim_{t \rightarrow \infty} x(t) = \alpha \mathbf{1}$. In other words, if there are some closed agents, all of whom agree on the state α , then the system converges to that consensus value.

- If $I_{closed} \neq \emptyset$, $\exists i, j \in I_{closed} : x_i \neq x_j$, and a permutation matrix P structures the adjacency matrix W of the network so that the closed agents I_{closed} precede the open agents I_{open} in PWP^T , then $\lim_{t \rightarrow \infty} x(t) = x^* = P^T[x_1^{*T}, x_2^{*T}]^T$, where x_1^* are the initial states of the closed agents, and $x_2^* = (I - W_{22})^{-1}W_{21}x_1^*$ if $|I_{open}| > 0$ and $x_2^* = []_{0 \times 1}$ otherwise. In other words, if there are multiple closed agents disagreeing on the state, then the system converges to the defined above state x^* of disagreement.

Proof. As the behavior of the system varies across the state space—as illustrated in Figure 3.3—let us, first, partition the latter and, then, prove convergence for each part individually.

$$[-1, 1]^n = \lim_{\epsilon \rightarrow +0} S_0(\epsilon) \cup \lim_{\epsilon \rightarrow +0} S_-(\epsilon) \cup \lim_{\epsilon \rightarrow +0} S_+(\epsilon) \cup \bigcup_{|I_{max} \cup I_{min}| \geq 2} S_*(I_{max}, I_{min}),$$

$$S_0(\epsilon) = [-1 + \epsilon, 1 - \epsilon]^n,$$

$$S_-(\epsilon) = \{Px \mid x \in \cup_{k=1}^n \{-1\}^k \times [-1, 1 - \epsilon]^{n-k}\},$$

$$S_+(\epsilon) = \{Px \mid x \in \cup_{k=1}^n \{1\}^k \times [-1 + \epsilon, 1]^{n-k}\},$$

$$S_*(I_{max}, I_{min}) = \{x \mid i \in I_{max} \leftrightarrow x_i = 1 \wedge i \in I_{min} \leftrightarrow x_i = -1\}.$$

Above, $S_0(\epsilon)$ corresponds to the interior of the state space, comprised of the states without extremist agents; $S_-(\epsilon)$ and $S_+(\epsilon)$ correspond to the parts of the state space's surface where all extremist agents are either in state -1 or in state 1 , respectively; and $S_*(I_{max}, I_{min})$ ($|I_{max} \cup I_{min}| \geq 2$) are the “edges” in which there are necessarily multiple extremist agents having different opinions. Now, we will study convergence of system (3.2) inside each of the above defined parts of the state space.

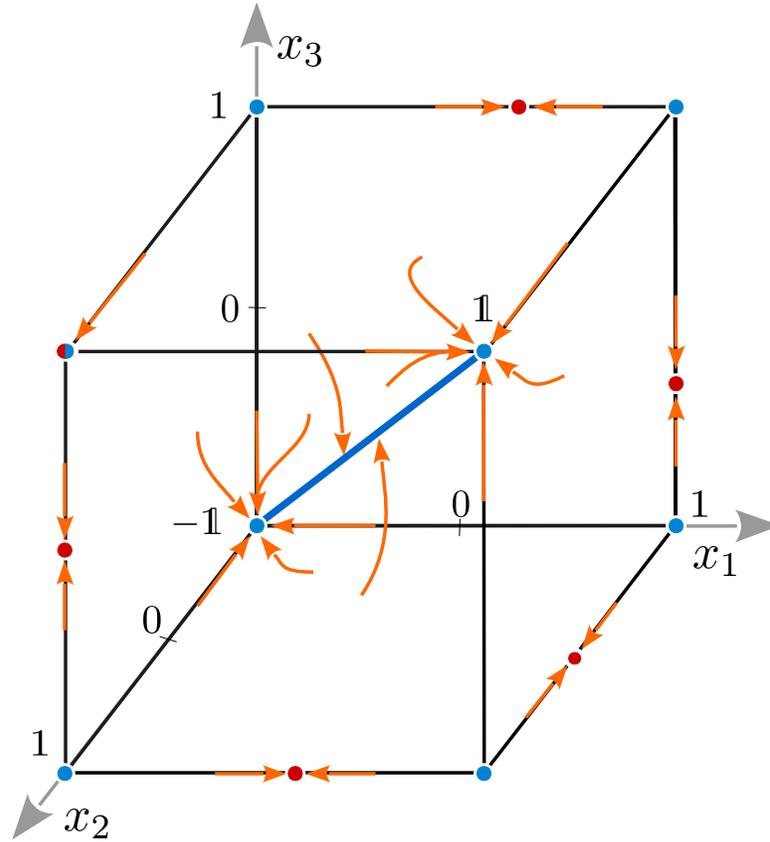


Figure 3.3: Convergence behavior of the model with stubborn extremists in three dimensions. A few representative trajectories are displayed as solid arrows; the solid diagonal correspond to the consensus equilibrium states of the system; the circles correspond to the disagreement equilibrium states; the corners are the states of disagreement where all agents are closed, while in the states in the interior of the cube’s edges, some agents remain open (depending on the location of closed agents in the network, some edges may have no such internal points of equilibrium, like in the case of edge $[-1,-1,1]-[-1,1,1]$).

(i) *Convergence in $S_0(\epsilon)$* (see Figure 3.4): From the general contraction Lemma 5, it follows that $S_0(\epsilon)$ is forward invariant with respect to (3.2). Additionally, inside $S_0(\epsilon)$, the agents cannot hold extreme opinions, so $A_{ii}(x) > 0$. Thus, it immediately follows from the general convergence Theorem 11 that all trajectories of system (3.2) starting in $S_0(\epsilon)$ converge to the latter’s subset of consensus states, that is, $N_0(\epsilon) = \{\alpha \mathbb{1} \mid \alpha \in [-1 + \epsilon, 1 - \epsilon]\}$, as $t \rightarrow \infty$.

(ii) *Convergence in $S_-(\epsilon)$* (see Figure 3.5): Forward invariance of $S_-(\epsilon)$ follows from the general contraction Lemma 5.

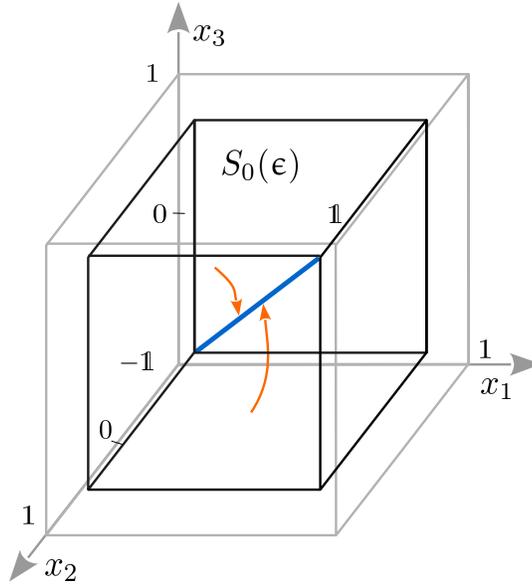


Figure 3.4: Convergence of the model with stubborn extremists in $S_0(\epsilon)$.

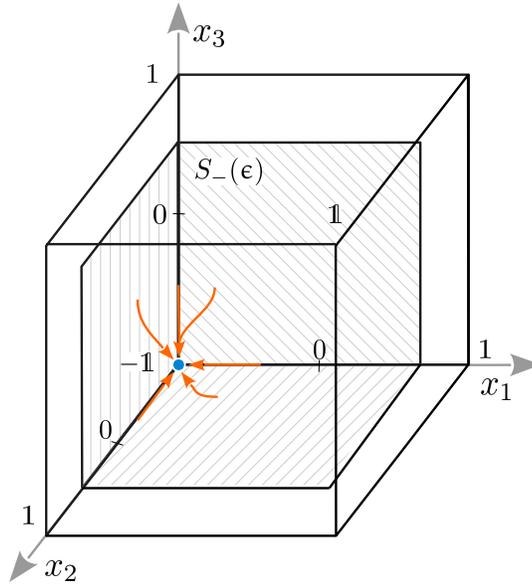


Figure 3.5: Convergence of the model with stubborn extremists in $S_-(\epsilon)$.

Additionally, since the states of $S_-(\epsilon)$ can only have extremists in state -1 , then $A_{ii}(x) = 0 \Leftrightarrow x_i = -1$, and, hence the zero-agreement property $\forall i, j \in 1, \dots, n : A_{ii}(x) = A_{jj}(x) = 0 \rightarrow x_i = x_j$ holds in $S_-(\epsilon)$. Thus, by Theorem 11, all trajectories of system (3.2) starting in $S_-(\epsilon)$ converge to $N_-(\epsilon) = S_-(\epsilon) \cap \{\alpha \mathbb{1} \mid \alpha \in [-1, 1]\} = \{-\mathbb{1}\}$ as $t \rightarrow \infty$.

(iii) *Convergence in $S_+(\epsilon)$* : The case is identical to the case of $S_-(\epsilon)$, with the extreme state 1 replacing the extreme state -1 , and the set to which convergence is expected to occur being $N_+(\epsilon) = \{\mathbb{1}\}$.

(iv) *Convergence in $S_*(I_{\max}, I_{\min})$* (see Figure 3.6): In this case, the general convergence Theorem 11 is not applicable, as we expect convergence to a state $x^* = P^\top[x_1^{*\top}, x_2^{*\top}]^\top$ of disagreement, with x_1^* corresponding to the initial states of the closed agents of $I_{\text{closed}} = I_{\max} \cup I_{\min} = \{i \mid |x_i| = 1\}$, $x_2^* = (I - W_{22})^{-1}W_{21}x_1^*$, and the adjacency matrix W being structured according to the partition $I_{\text{closed}}, I_{\text{open}}$.

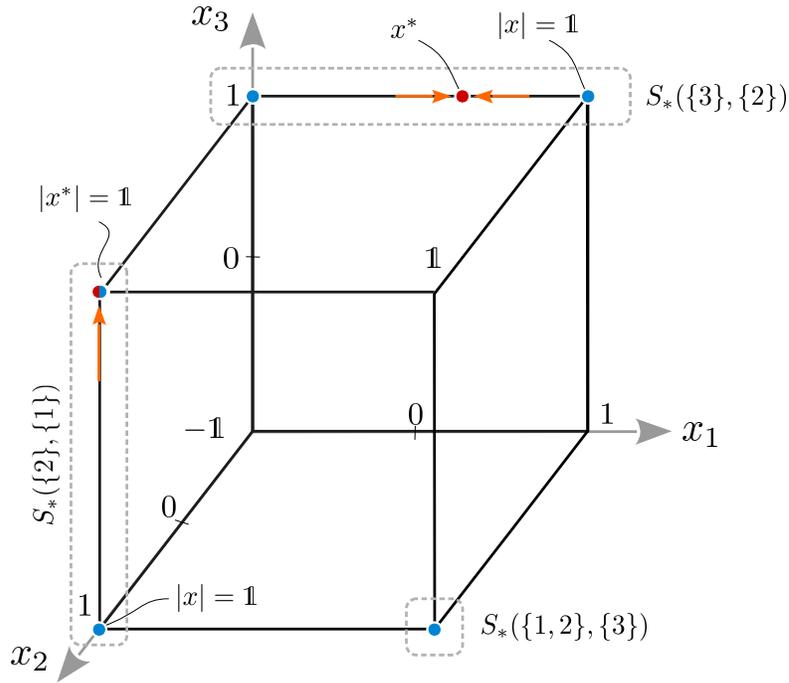


Figure 3.6: Convergence of the model with stubborn extremists in $S_*(I_{\max}, I_{\min})$.

Notice, that if either $|I_{\min} \cup I_{\max}| = n$, or if the open agents can be reached only from the extremists in one state, then $|x^*| = \mathbb{1}$, that is, the states of all the agents may asymptotically become extreme.

First, we will construct a Lyapunov function out of max-min functions, then, re-use it to prove invariance and, eventually, convergence to a state of disagreement using the Invariance

Principle.

Consider function $V_{*\max}(x) = \max(x - x^*)$, where x^* is the state of disagreement defined above. From Theorem 7, it follows that

$$\partial V_{*\max}(x) = \{P^\top[\alpha^\top, \mathbb{0}^\top]^\top\},$$

where α , consisting of the coefficients of a convex combination, correspond to the agents of $I_{*\max}(x) = \{i \mid (x - x^*)_i = \max(x - x^*)\}$, $\mathbb{0}$ correspond to the rest of the agents, and P^\top is a permutation matrix restoring the original agent order. Now, for any $x \in S_*(I_{\max}, I_{\min})$ and $\xi \in \partial V_{*\max}(x)$, up to reordering of the agents, we will have

$$\langle \xi, f(x) \rangle = -\alpha^\top A_{*\max}(x)(x_{*\max} - [W_{11}, W_{12}]x), \quad (3.6)$$

where $A_{*\max}(x)$ are the susceptibilities of the agents of $I_{*\max}(x)$, and the adjacency matrix is structured according to the agent set partition $I_{*\max}(x)$, $\{1, \dots, n\} \setminus I_{*\max}(x)$. Our immediate goal is to determine the sign of the obtained expression (3.6). To that end, consider factor $(x_{*\max} - [W_{11}, W_{12}]x)$, letting $x_{*\max}^*$ be the part of x^* corresponding to the agents of $I_{*\max}(x)$

$$\begin{aligned} x_{*\max} - [W_{11}, W_{12}]x &= x_{*\max} - \underbrace{x_{*\max}^* + x_{*\max}^*}_{\mathbb{0}} - [W_{11}, W_{12}]\underbrace{(x - x^* + x^*)}_{\mathbb{0}} \\ &= (x_{*\max} - x_{*\max}^*) - [W_{11}, W_{12}](x - x^*) + \underbrace{(x_{*\max}^* - [W_{11}, W_{12}]x^*)}_{\mathbb{0}} \\ &= (x - x^*)_{*\max} - [W_{11}, W_{12}](x - x^*). \end{aligned}$$

It is clear from the definition of $I_{*\max}(x)$ that $(x - x^*)_{*\max} - [W_{11}, W_{12}](x - x^*) \geq 0$ and, hence $\langle \xi, f(x) \rangle \leq 0$. Furthermore, we can apply the argument from the proof of the general convergence Theorem 11, to establish that when there is an agent i such that $((x - x^*)_{*\max} -$

$[W_{11}, W_{12}](x - x^*)_i = 0$, we can vary α to make $\langle \xi, f(x) \rangle$ take different values for the same x . Thus, we can conclude that $\widetilde{\mathcal{L}}_f V_{*\max}(x) = \{0\}$ when $|x| = 1$ (as $\xi = 0$) or when $x = x^*$, and $\max \widetilde{\mathcal{L}}_f V_{*\max}(x) < 0$ for the other $x \in S_*(I_{\max}, I_{\min})$.

We can repeat the same reasoning to establish that, for $V_{-*\min} = -\min(x - x^*)$ it holds that $\widetilde{\mathcal{L}}_f V_{-*\min}(x) = \{0\}$ when $|x| = 1$ or $x = x^*$, and $\max \widetilde{\mathcal{L}}_f V_{-*\min}(x) < 0$ for the rest of $x \in S_*(I_{\max}, I_{\min})$.

Our reasoning about $V_{*\max}(x)$ and $V_{-*\min}(x)$ allow us to conclude that function

$$V_{*max-min}(x) = V_{*\max}(x) + V_{-*\min}(x) = \max(x - x^*) - \min(x - x^*)$$

is a Lyapunov function for system (3.2), as required by the Invariance Principle. Additionally, $S_*(I_{\max}, I_{\min})$ is forward invariant, which immediately follows from the agent subset invariance Lemma 4. Thus, by Invariance Principle, all trajectories of system (3.2) starting in $S_*(I_{\max}, I_{\min})$ converge to set

$$N_*(I_{\max}, I_{\min}) = \{x \mid |x| = 1\} \cup \{x^*\},$$

in which $0 \in \widetilde{\mathcal{L}}_f V_{*max-min}(x)$. What remains to show is what element of $N_*(I_{\max}, I_{\min})$ the system converges to.

Clearly, if all the agents are initially closed, that is, $|x(0)| = 1$, then $\lim_{t \rightarrow \infty} x(t) = x(0)$, which follows from the agent subset invariance Lemma 4. Now, assume that $|x(0)| \neq 1$. In such a case, a trajectory cannot approach any element of $\{x \mid |x| = 1\}$ (except, possibly, x^* in the case when the open agents are only reachable by the closed agents having the same state, and, as a result, $|x^*| = 1$), as, generally, approaching one of these states would violate at least

one of the above proven inequalities

$$\begin{aligned}\max \widetilde{\mathcal{L}}_f V_{*\max}(x) &= \max \widetilde{\mathcal{L}}_f \max(x - x^*) \leq 0, \\ \max \widetilde{\mathcal{L}}_f V_{*\min}(x) &= \max \widetilde{\mathcal{L}}_f (-\min(x - x^*)) \leq 0.\end{aligned}$$

Hence, if $|x(0)| \neq 1$, then the trajectories of (3.2) converge to x^* as $t \rightarrow \infty$. ■

3.6 Simulation Results

In addition to the theoretical results of Section 3.5, we report simulation results for each of the proposed models over two synthetic and one real-world networks: Erdős-Rényi ($|V| = 200$, $|E| = 10326$, $\mathbb{P}_{edge} = .25$), Scale-Free ($|V| = 200$, $|E| = 944$, $\gamma = -2.3$), and Zachary's Karate Club [164] ($|V| = 34$, $|E| = 156$). The edge weights were selected uniformly at random. For each network and each model, we consider several qualitatively different cases of a randomly chosen initial state $x(0)$. We run simulation either until convergence, or, if convergence is slow, long enough, so that the fact of convergence and the limiting state both become evident. We display evolution of each component of the solution. The figures are shown on Page 104.

3.7 Discussion

In this section, we summarize and interpret the results obtained in the Chapter, as well as assess where they fit in and how contribute to the existing body of research.

New models: We have defined the general model of polar opinion dynamics

$$\dot{x} = -A(x)Lx, \tag{3.1}$$

that, depending on a specific form of the susceptibility functions on the main diagonal of $A(x)$, has interpretation in terms of one of the socio-psychological theories. Model (3.1) can be viewed as a non-linear analog of Abelson [49], DeGroot [91], and Friedkin-Johnsen [98] models, with the dependence of the agents' susceptibilities $A(x)$ to persuasion upon their current opinions being the key distinguishing trait of our model. Mathematically, model (3.1) is also related to the class of bounded confidence models [131, 121, 122, 123, 124], in which the opinion-adoption behavior of the agents also depends on the agent's current beliefs, yet, this dependence is based upon the socio-psychological principles different from ours. A notable exception is the work of Sobkowicz [131], in which, the author uses the opinion resilience mechanisms similar to the ones we use in our models with stubborn positives and stubborn extremists.

Behavior of the general model: For the general model (3.1), in Theorem 11, we have provided a sufficient condition for the convergence to consensus. Roughly speaking, a trajectory starting inside a forward invariant set approaches a state of consensus if all closed agents have similar states, that is, $A_{ii}(x) = A_{jj}(x) = 0 \rightarrow x_i = x_j$. From the sociological perspective, it means that, as long as all the ultimately stubborn agents in the network agree upon their states, the agents will eventually agree, as there is no force that would drive the system to disagreement. The observed behavior is different from Friedkin-Johnsen model in that, in our model, the presence of non-fully open agents, having $A_{ii}(x) < 1$, does not immediately lead to an asymptotic disagreement. The obtained sufficient condition for convergence to consensus, compared to its analogs derived in [127, 125], is better interpretable from the sociological point of view, and is not more restrictive.

Behavior of the specialized models: In addition to the general model, we have considered three specialized models

$$\dot{x} = -\frac{1}{2}(I - \text{diag}(x))Lx, \quad (\textit{stubborn positives})$$

$$\dot{x} = -\text{diag}(x)^2Lx, \quad (\textit{stubborn neutrals})$$

$$\dot{x} = -(I - \text{diag}(x)^2)Lx. \quad (\textit{stubborn extremists})$$

The behavior of these models, studied in Theorems 12, 13, 14, can be summarized as follows.

Convergence to consensus (closed agents absent): If there are no closed agents in the network, that is, if $A(x(0)) > 0$, and, as a result, every agent is at least to some degree susceptible to persuasion, then the system converges to a state of consensus, as shown for the example of the model with stubborn positives in Figure 3.7.

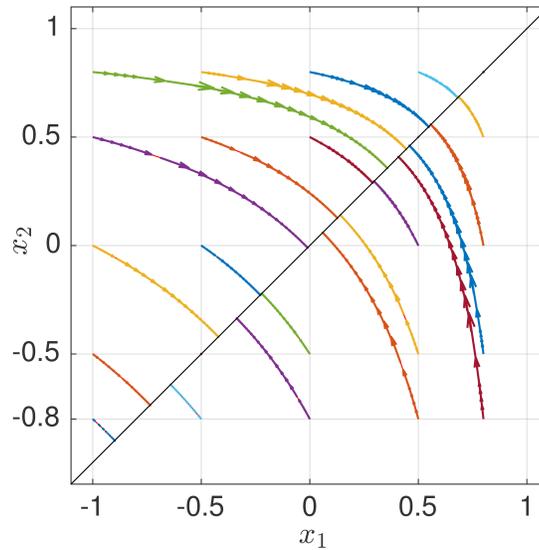


Figure 3.7: Phase portrait: convergence to consensus of the model with stubborn positives in the absence of closed agents.

In the absence of closed agents, the particular consensus value is known only for the model with stubborn neutrals when the agents with both positive and negative states are present—in this case, the system converges to 0 .

This behavior is not surprising—when term $A(x)$ of the vector field $-A(x)Lx$ does not prevent any agent from changing its state, the negative Laplacian expectedly drives the state toward a consensus. Thus, if there are no ultimately stubborn agents, then the group will asymptotically reach an agreement.

Convergence to consensus (closed agents present): In the presence of closed agents all of whom agree on their state $\alpha \in [-1, 1]$, the system converges to consensus $\alpha\mathbb{1}$. In other words, if the ultimately stubborn agents are present and share the same opinion, they will persuade the rest of the group to adopt that opinion. A representative example of such behavior is given in Figure 3.8, for the model with stubborn neutrals.

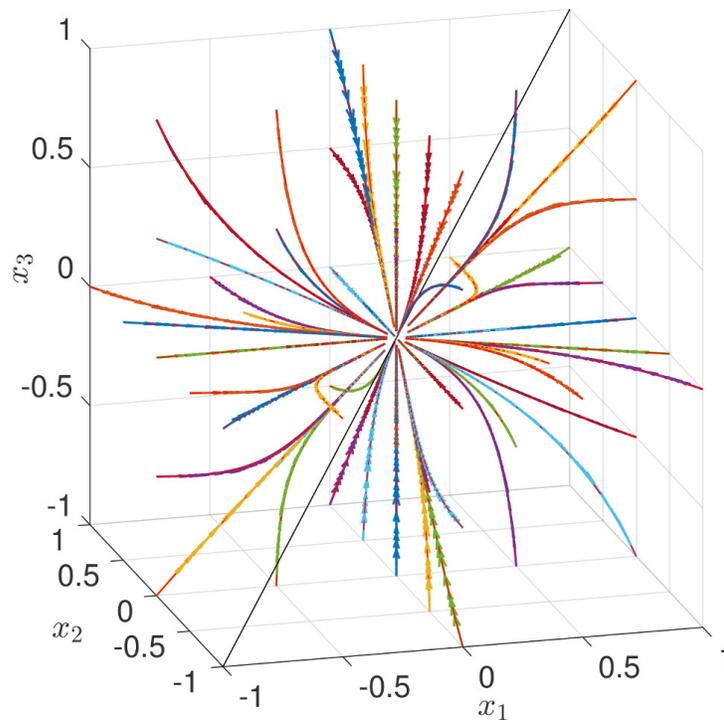


Figure 3.8: Phase portrait: convergence to consensus of the model with stubborn neutrals when closed agents are present.

A natural conclusion is that the only force that can counteract the persuasion efforts of the ultimately stubborn agents agreeing on an opinion is the ultimately stubborn agents having a different opinion.

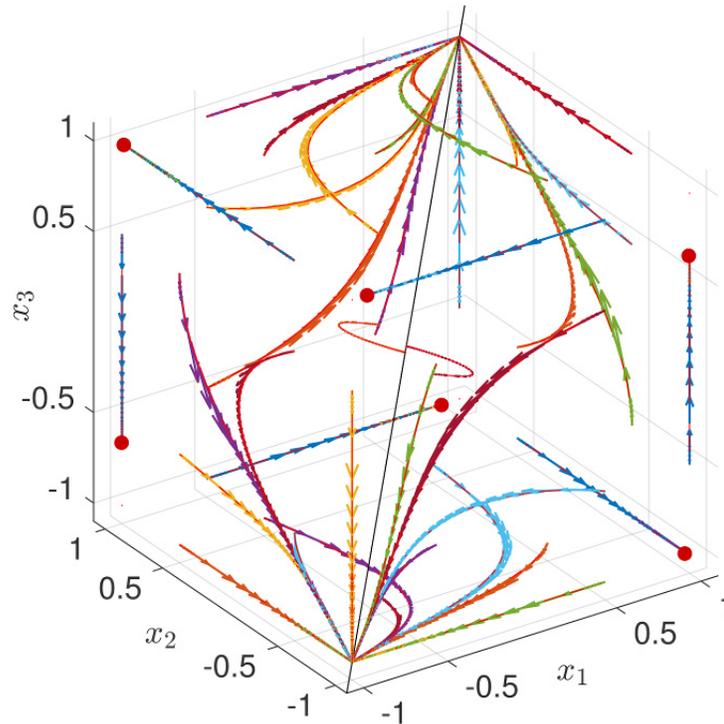


Figure 3.9: Phase portrait: asymptotic behavior of the model with stubborn extremists. Strictly inside the state space, trajectories converge to a consensus. On the surface of the state space, when all closed agents have the same state α , trajectories converge to that consensus state $\alpha\mathbf{1}$. In the presence of multiple closed agents holding different opinions, the system converges to a state of disagreement (such states are displayed as solid circles).

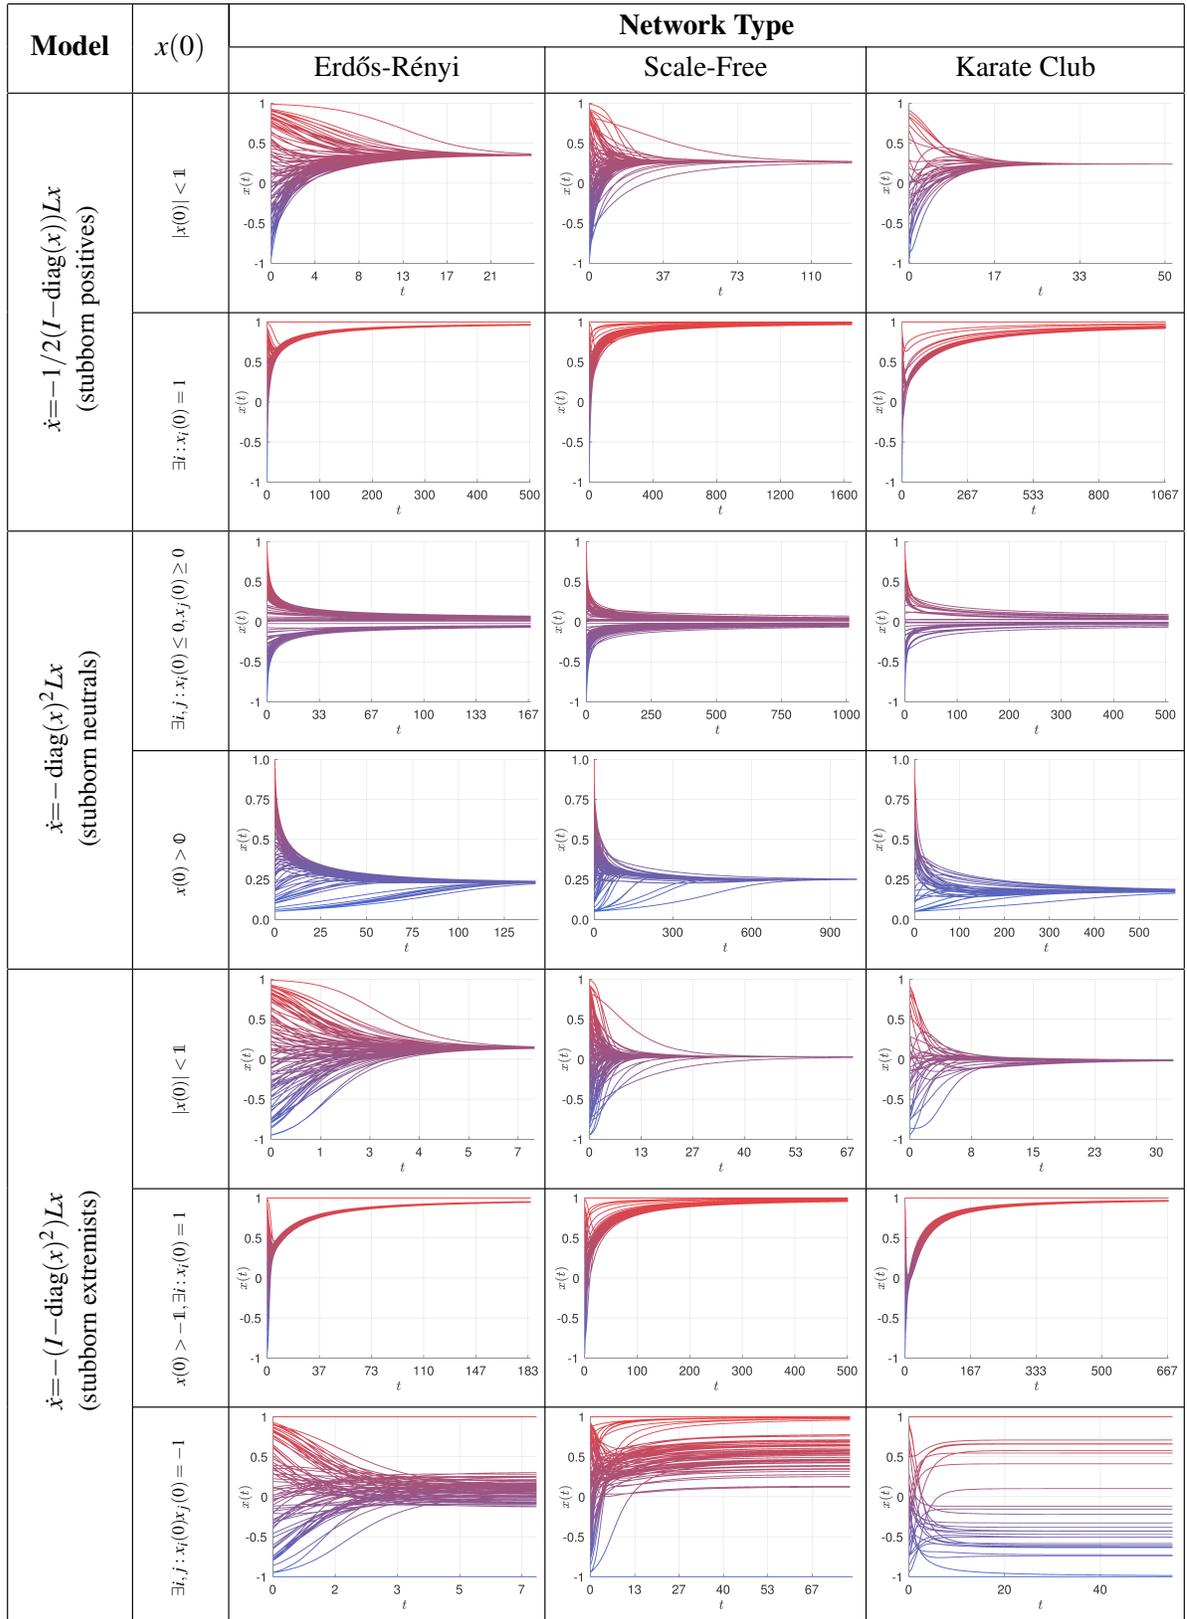
Convergence to disagreement: Finally, in the presence of multiple closed agents holding different opinions, which is possible only for the model with stubborn extremists, the system converges to a state of disagreement. Figure 3.9 shows a full range of qualitatively different asymptotic behaviors of the model with stubborn extremists. If, in addition to the closed agents holding different opinions, there are some open agents in the network, then the closed agents will persuade the open agents to adopt a combination of their opinions. In this case, the particular limiting state to which the system will converge will depend on the structure of the network and the locations and states of the closed agents, yet, not on the initial opinions of the open agents. (A similar behavior has also been observed in the context of the Voter model with stubborn agents [155, Sec. 4]). It is particularly interesting that the opinions of the open

agents in the latter limiting state have a rather simple expression $(I - W_{22})^{-1}W_{21}x_1^*$, given in Theorem 14, where W_{22} is the block of the adjacency matrix corresponding to the cluster of the open agents, W_{21} is the block responsible for the influence of the closed agents upon the open agents, and x_1^* are the closed agents' (initial) states.

Model analysis: The bulk of our theoretical analysis of the models' behavior is comprised of the proofs of convergence. The standard tools for the analysis of convergence of non-linear models, such as LaSalle Invariance Principle, require existence of a smooth Lyapunov function, with quadratic functions being a popular choice. The latter, however, may be hard and, sometimes, provenly impossible [113] to find for a model defined over a general directed network. In this chapter, we show, using several existing tools from non-smooth analysis, how to apply non-smooth max-min functions to prove convergence of our models. Such Lyapunov functions have been considered in the literature [113, 165], however, this work is the first to provide a full formal analysis of such functions used along with the generalized Invariance Principle. Due to the generality of the non-smooth analysis tools we have used, our analysis can be easily adapted to other non-linear models defined over directed networks, with Lyapunov functions constructed out of convex components.

Acknowledgement

This chapter—largely based upon the author's published work [166]—has benefited from advice of multiple people. We express our gratitude to Noah E. Friedkin for pointing out the usefulness of having multiple definitions of the agent susceptibility, providing critique of an early version of the manuscript, as well as introducing us to the field of social psychology. We are also grateful to Andrew R. Teel for a useful discussion. Finally, we thank the anonymous reviewers of IEEE TRANSACTIONS ON AUTOMATIC CONTROL for their contribution.



Chapter 4

Defense Against Social Control via Link Recommendation

Existing socio-psychological studies show that the process of opinion formation is inherently a network process, and user opinions in a social network are attracted to a certain average belief of the entire network. One simple and intuitive incarnation of this notion of the opinion attractor is the average $\pi^\top x$ of user opinions x_i weighted by the users' eigenvector centralities π_i . This value is also known as the *asymptotic consensus value* due to its being the asymptotic limit to which user opinions converge under DeGroot opinion formation model. It is clear that the opinion attractor is a lucrative target for control, as altering it essentially changes the average belief in the network. Since any potentially malicious control of the opinion distribution in a social network is undesirable, it is important to design methods to prevent external attacks upon the asymptotic consensus value.

This chapter is dedicated to the design of one mechanism to defend a network against social control. We assume that there is an adversary that aims to maliciously change the asymptotic consensus value by altering the opinions of some users in the network. We, then, state DIVER—an NP-hard problem of disabling such external attacks upon the user opinion distri-

bution via strategically altering the network’s eigenvector centrality by the means of recommending a limited number of links to the network’s users. Relying on the theory of Markov chains, we provide perturbation analysis that shows how eigencentality and, hence, DIVER’s objective function change in response to an edge’s addition to the network. The latter leads to the design of a pseudo-linear-time heuristic for DIVER, relying on efficient estimation of mean first passage times in a Markov chain. We confirm our theoretical and algorithmic findings, and study effectiveness and efficiency of our heuristic in experiments with synthetic and real-world networks.

This chapter is organized as follows. We, first, motivate the problem in Section 4.1, and, having defined preliminaries in Section 4.2, formally state our problem and prove its hardness in Section 4.3. In Section 4.4, we review related work on analytic and combinatorial optimization of network topology as well as centrality perturbation. Our main result on optimal strategic link recommendation is developed in Section 4.5. We buttress our theoretical findings with experiments on synthetic and real-world networks in Section 4.6, and conclude with the discussion of our results and future work in Sections 4.7 and 4.8, respectively.

4.1 Introduction

It is well-known that in the absence of the objective means for opinion evaluation, people tend to evaluate their opinions by comparison with the opinions of others. Thus, social networks impact the opinion formation process in the society. It is clearly desirable that this process is natural and fair, with good ideas emerging and spreading, and bad ideas declining and disappearing. However, viral marketing experts may be interested in affecting the opinion formation process with the goal of driving the opinion distribution to a certain business-imposed objective. One popular way of affecting—or controlling—opinion formation is influence maximization, whose central idea is to affect the opinions of some network users with the goal

of maximizing the subsequent spread of “right” opinions from these users throughout the network, or, more generally, shifting the opinion distribution towards a desired state. Naturally, the society would benefit from a mechanism preventing such potentially malicious interventions into the opinion formation process in social networks. This chapter is dedicated to the design of one such mechanism—a link recommendation algorithm that disables the effect of the attempts to control the opinion distribution in a social network.

The notion central to our work is that of the average opinion of a social network, that characterizes the entire network’s general belief. For example, it can measure to what extent, on average, network users prefer one smartphone brand over the other. Based on the well-established socio-psychological theories, such as Festinger’s social comparison [90] and cognitive dissonance [92] theories, the opinions of users in the network are attracted towards the average opinion. However, the basic arithmetic average would not suffice, as the opinions of more “central” users—such as celebrities—are expected to contribute more to the entire network’s belief.

In order to formalize the notion of the average opinion, we recall DeGroot opinion formation model [91] rooted in the above mentioned socio-psychological theories:

$$x(t+1) = Wx(t),$$

$$x(t) \in [0, 1]^n, W \in [0, 1]^{n \times n}, W\mathbf{1} = \mathbf{1},$$

where $x(t) \in [0, 1]^n$ is a vector of quantified opinions of all users at time t , and $W \in [0, 1]^{n \times n}$ is a row-stochastic interpersonal appraisal matrix—playing the role of the adjacency matrix of a directed social network—whose element w_{ij} quantifies the relative extent to which user i values the opinion of user j (see Figure 4.1a). According to this model, users form their opinions via weighted averaging of their own opinions with those of their neighbors in the network. It is well-known that, under DeGroot model, in a “well-connected” network, the opinions of all

users approach the same value—the *asymptotic consensus value*

$$\lim_{t \rightarrow \infty} x_i(t) = \langle \pi, x(0) \rangle = \pi^\top x(0) = \sum_{j \in V} \pi_j x_j(0),$$

being the sum of the users’ initial opinions $x_j(0)$, weighted by the users’ eigenvector centralities π_j . While in real-world situations, and contrary to what DeGroot model prescribes, people do not always arrive at the same opinion, the opinions’ attraction to the (weighted) average opinion is exactly what socio-psychological theories postulate. Thus, we adopt the asymptotic consensus value $\langle \pi, x \rangle$ as the formalization of the notion of the network’s average opinion, dropping the time notation for readability.

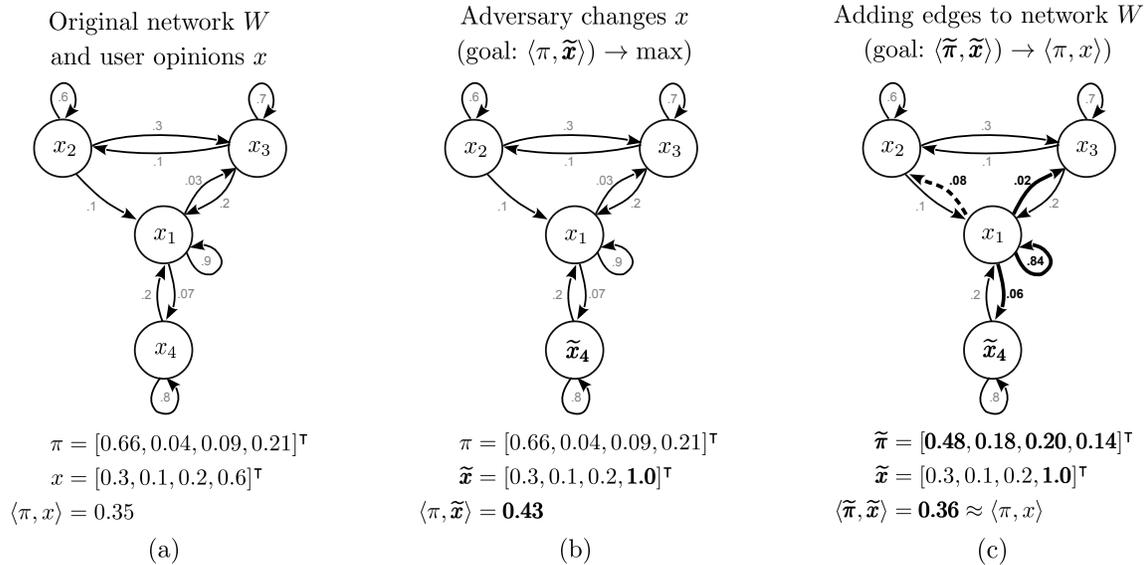


Figure 4.1: DIVER’s definition by example. The adversary influences user opinions, $x \rightarrow \tilde{x}$, increasing the asymptotic consensus value $\langle \pi, x \rangle \rightarrow \langle \pi, \tilde{x} \rangle$. The network responds by adding edges, reducing the asymptotic consensus value $\langle \tilde{\pi}, \tilde{x} \rangle$, driving it back to original $\langle \pi, x \rangle$.

Since the asymptotic consensus value plays the role of an opinion attractor, it is a lucrative target for influence. We assume that there is an external adversary whose goal is, without loss of generality, to maximize the asymptotic consensus value $\langle \pi, x \rangle$. To that end, the adversary influences some users in the network, changing the opinion distribution $x \rightarrow \tilde{x}$ and, thus, the

asymptotic consensus value $\langle \pi, x \rangle \rightarrow \langle \pi, \tilde{x} \rangle$, as illustrated in Figure 4.1b. Our goal is to respond to this attack, and restore the asymptotic consensus value to its original state $\langle \pi, x \rangle$. We, however, cannot directly influence user opinions; the only legitimate opinion control tool available to us is *edge recommendation*. We add a limited number of edges, thereby, changing eigencentralities $\pi \rightarrow \tilde{\pi}$ and restoring the original asymptotic consensus value, $\langle \pi, \tilde{x} \rangle \rightarrow \langle \tilde{\pi}, \tilde{x} \rangle \approx \langle \pi, x \rangle$, as illustrated in Figure 4.1c.

Our *central goal* in this chapter is to design a scalable algorithm that—under the above described attack upon the opinion distribution—would identify the edges whose addition to the social network would efficiently drive the altered asymptotic consensus value back to its state prior to the attack, nullifying the attack’s impact. Our *specific contributions* are:

- We define DIVER—a *new problem* of disabling external influence in a social network via edge recommendation—and prove its NP-hardness.
- We provide a *novel perturbation analysis*, establishing how the network’s eigencentralities vector changes in response to edge addition to the network. This analysis leads to the definition of an edge score $f_\pi(i, j)$ quantifying the potential impact of the addition of directed edge (i, j) to the network upon DIVER’s objective.
- We show how to *estimate edge scores* f_π in *pseudo-constant time* in networks with skewed eigenvector centrality distribution, such as scale-free networks.
- We provide a *pseudo-linear-time heuristic* for DIVER relying on edge scores f_π , and demonstrate its effectiveness and efficiency in experiments with synthetic and real-world networks.

4.2 Preliminaries

We are given a sparse directed strongly connected aperiodic social network $G(V, E)$, $|V| = n$, $|E| = O(n)$, having row-stochastic adjacency matrix $W \in [0, 1]^{n \times n}$, $W\mathbf{1} = \mathbf{1}$ —also known as the interpersonal appraisal matrix—whose entry $w_{ij} \in [0, 1]$ reflects the relative extent to which user i takes into account the opinion of user j while forming his or her opinion. Aperiodicity can be replaced by the requirement of the network’s having at least one self-loop with a non-zero weight [167], which translates into a natural requirement of having at least one user who does not completely disregard his or her own opinion in the process of new opinion formation.

The average belief of the network—or, alternatively, the value to which the opinions of the users are being attracted—is represented by the *asymptotic consensus value* $\langle \pi, x \rangle$, where $x \in [0, 1]^n$ are opinions of the users, and $\pi \in \mathbb{R}_+^n$, $\|\pi\|_1 = 1$ is the ℓ_1 -normalized dominant left eigenvector of W . From its definition, π is also the vector of eigencentralities of the network’s nodes, and can also be viewed as the vector of no-teleportation PageRank scores, or the stationary distribution of the ergodic Markov chain with state transition matrix W . Due to the latter, we may refer to W as a Markov chain, and rely upon the following characteristic of W if viewed as a chain.

Definition 7 (Mean First Passage Time). *The mean first passage time (MFPT) m_{ij} from state i to state j of a Markov chain is the expected number of steps it takes the chain started in state i to pass through state j . m_{ii} is the mean first return time (MFRT) to state i .*

The following Theorems 15 and 16 immediately follow from Theorems 4.4.4 and 4.4.5 of Kemeny and Snell [168], respectively, with the regularity of Markov chain in [168] being translated into our aperiodicity and strong connectivity of the chain’s network.

Theorem 15 (Connection Between MFRT and π). *For any state i of Markov chain W with an aperiodic strongly connected network, $m_{ii} = 1/\pi_i$.*

Theorem 16 (MFPT One-Hop Conditioning). *For any states i and j of Markov chain W with an aperiodic strongly connected network, $m_{ij} = 1 + \sum_{k \neq j} w_{ik} m_{kj}$.*

We summarize our notation in Table 4.1.

$\mathbf{1}$	vector of all ones
$\text{diag}(v)$	diagonal matrix with vector v at the main diagonal
e_i	i 'th column of the identity matrix
$x(\tilde{x})$	user opinions (user opinions altered by the adversary)
W	network's row-stochastic adjacency matrix
\tilde{W}	altered network' row-stochastic adjacency matrix
θ_{ij}	weight of added directed edge (i, j)
\mathbb{P}_{ij}	acceptance probability of added directed edge (i, j)
$\pi(\tilde{\pi})$	ℓ_1 -normalized left dominant eigenvector of W (of \tilde{W})
m_{ij}	mean first passage time from i to j in chain W

Table 4.1: Summary of notation.

4.3 Problem's Statement and Hardness

Given a directed social network, at each point in time, we can observe its users' opinions. We assume that at some point the external adversary makes an influence maximization attempt by targeting several users and changing their opinions, with the goal of, without loss of generality, maximizing the asymptotic consensus value $\langle \pi, x \rangle$. We assume no knowledge of which users have been attacked, relying only on our ability to detect such attacks using existing methods that range from the anomalous event detection in a scalar time series of average opinion values [169], to tracking whether the current cumulative changes in user opinions follow a pattern prescribed by the solution to an influence maximization problem, to network sensing for outbreak detection [170], to opinion dynamics model-driven anomaly detection techniques [171] from Section 2.7 of Chapter 2 of this thesis.

Having detected an external influence attempt, we are given the opinion distribution $x \in [0, 1]^n$ preceding the attack as well as the externally altered opinion distribution $\tilde{x} \in [0, 1]^n$. As

a result of the attack, the original asymptotic consensus value $\langle \pi, x \rangle$ has changed to $\langle \pi, \tilde{x} \rangle$. Our goal is to add a limited number k of edges to the network and, thereby, change π in such a way, that the resulting asymptotic consensus value $\langle \tilde{\pi}, \tilde{x} \rangle$ is as close as possible to its state $\langle \pi, x \rangle$ before the attack. Formally, the problem of disabling external influence via edge recommendation is defined as follows:

$$\boxed{\text{DIVER}(W, k, x, \tilde{x}) = \arg \min_{\tilde{W}} |\langle \tilde{\pi}(\tilde{W}), \tilde{x} \rangle - \langle \pi, x \rangle|}, \quad (4.1)$$

where the perturbed row-stochastic adjacency matrix \tilde{W} differs from W by k new edges whose weights θ_{ij} we cannot control—since they correspond to interpersonal appraisals that users choose themselves—yet, can estimate using existing techniques [53], and, hence, assume the knowledge of.

Notice that, after new edge addition to a user's out-neighborhood, the existing edges are proportionally downweighted, and \tilde{W} is kept row-stochastic. For $k = 1$, after addition of directed edge (r, c) with weight θ_{rc} , the perturbed adjacency matrix \tilde{W} is defined as follows

$$\tilde{W} = W - \theta_{rc} \text{diag}(e_r)W + \theta_{rc} e_r e_c^\top. \quad (4.2)$$

The same single-edge perturbation is illustrated in Figure 4.2.

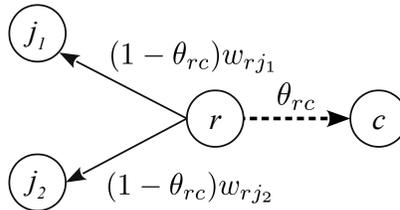


Figure 4.2: Perturbation of a network with a single edge (r, c) with weight θ_{rc} . Existing edges are downweighted, maintaining row-stochasticity of the adjacency matrix.

We, first, focus on deterministic edge addition and provide a straightforward extension of our results with edge acceptance probabilities in Section 4.5.6.

DIVER's complexity comes along two dimensions—searching for the best edge subset delivering the minimum of the objective, and assessing the impact of a given subset of edges upon the objective. While the latter can be done in polynomial time, the edge subset search cannot and is the root cause of NP-hardness. In the following Theorem 17, we establish that, even for the case of undirected networks, $\text{DIVER}(W, k, x, \tilde{x})$ is NP-hard.

Theorem 17 (Hardness of DIVER). *$\text{DIVER}(W, k, x, \tilde{x})$ is NP-hard for undirected networks.*

Proof. In the proof, we will show that

$$\text{DIVER}(W, k, x, \tilde{x}) = \arg \min_{\tilde{W}} |\langle \tilde{\pi}(\tilde{W}), \tilde{x} \rangle - \langle \pi, x \rangle|$$

applied to a certain simple undirected network can be used as a solver for the classic NP-complete subset sum problem. For readability, we will abuse notation and assume that the value of $\text{DIVER}(W, k, x, \tilde{x})$ is the minimum value itself, rather than the corresponding arg min. The remainder of the proof will consist of four parts.

1) *Subset sum problems:* The subset sum problem $\text{SSP}(\{z_i\}, s)$ is a classic NP-complete problem of deciding whether a given finite set $\{z_i\} \subset \mathbb{Z}^n$ of integers has a non-empty subset with a predefined sum $s \in \mathbb{Z}$. (SSP appears on Karp's list of NP-complete problems [172, p.95] under the name KNAPSACK). A related problem is the problem $\text{kSSP}_{01}(\{z_i\}, k, s)$ of deciding whether, among a finite number of bounded reals $z_i \in [0, 1]$, there is a non-empty subset of k

elements summing up to a given value $s \in [0, 1]$. Reduction $\text{SSP} \propto \text{kSSP}_{01}$ is as follows:

$$\begin{aligned} \text{SSP}(\{z_i\}, s) &= \bigvee_{k=1}^n \text{kSSP}_{01}(\{z_i''\}, k, s''), \\ z_i' &= z_i + L \in \mathbb{Z}_+, \quad s' = s + kL \in \mathbb{Z}_+, \\ L &= |\min\{0, \min\{s, \min z_i\}\}| \\ z_i'' &= z_i'/M \in [0, 1], \quad s_i'' = (s + kL)/M \in [0, 1], \\ M &= \max\{s', \max z_i'\}. \end{aligned}$$

2) *Undirected uniformly weighted networks and their eigenvector centrality*: Let $W_{01} \in \{0, 1\}^{n \times n}$ be the binary adjacency matrix of an undirected network, $d = W_{01}\mathbb{1}$ be a vector of node degrees, and $D = \text{diag}(d)$. Further, let $W = D^{-1}W_{01}$. We say that W is the adjacency matrix of an *undirected uniformly weighted network*, since all the edges within the same neighborhood are weighted equally. Notice that W is row-stochastic, as

$$W\mathbb{1} = D^{-1}W_{01}\mathbb{1} = D^{-1}d = \mathbb{1}.$$

Since $d^\top W = d^\top D^{-1}W_{01} = \mathbb{1}^\top W_{01} = d^\top$, vector $\pi = d/\|d\|_1 = d/(2|E|) = d/(2m)$ is the ℓ_1 -normalized dominant left eigenvector—or, eigenvector centrality—of W . If the underlying unweighted network W_{01} is perturbed with k undirected edges $(i, j) \in S$, $|S| = k$, then the eigenvector centrality of the corresponding weighted network becomes

$$\tilde{\pi} = \frac{1}{2(m+k)} \left(d + \sum_{(i,j) \in S} (e_i + e_j) \right) = \frac{1}{m+k} \left(m\pi + \sum_{(i,j) \in S} (e_i + e_j)/2 \right), \quad (4.3)$$

where e_i is the i 'th column of the identity matrix.

3) *DIVER in undirected uniformly weighted networks*: If network W is undirected uniformly weighted and, thus, defined by its binary adjacency matrix W_{01} , then DIVER's objective

function over such W can be rewritten as follows:

$$\begin{aligned}
f(\tilde{W}_{01}) &= |\langle \tilde{\pi}, \tilde{x} \rangle - \langle \pi, x \rangle| = (\text{from (4.3)}) = \\
&= \left| \frac{m}{m+k} \langle \pi, \tilde{x} \rangle + \frac{1}{2(m+k)} \left\langle \sum_{(i,j) \in S} (e_i + e_j), \tilde{x} \right\rangle - \langle \pi, x \rangle \right| \\
&= \frac{1}{m+k} \left| \left\langle \sum_{(i,j) \in S} (e_i + e_j), \tilde{x}/2 \right\rangle - \langle \pi, (m+k)x - m\tilde{x} \rangle \right| \\
&= a(k) \left| \left\langle \sum_{(i,j) \in S} (e_i + e_j), \tilde{x}/2 \right\rangle - b(k, x, \tilde{x}) \right|,
\end{aligned}$$

where $a(k) = (m+k)^{-1}$ and $b(k, x, \tilde{x}) = \langle \pi, (m+k)x - m\tilde{x} \rangle$. Since k , and, consequently, $a(k)$ are constant, minimization of $f(\tilde{W}_{01})$ is equivalent to minimization of

$$f'(\tilde{W}_{01}) = \left| \left\langle \sum_{(i,j) \in S} (e_i + e_j), \tilde{x}/2 \right\rangle - b(k, x, \tilde{x}) \right|. \quad (4.4)$$

4) *Reduction* $\text{kSSP}_{01} \propto \text{DIVER}$: Suppose we are given an instance $\text{kSSP}_{01}(z, k, s)$, with $z \in [0, 1]^n$, $k \in \mathbb{N}$, and $s \in [0, 1]$. In what follows, we will show that the solution to $\text{kSSP}_{01}(z, k, s)$ is obtained by checking whether

$$\min_{\tilde{W}^{KC}} \text{DIVER} \left(W^{KC}, k, \frac{s\mathbb{1} + m(z \otimes \mathbb{1}_2)}{m+k}, z \otimes \mathbb{1}_2 \right) = 0, \quad (4.5)$$

where W^{KC} is the adjacency matrix of an undirected uniformly weighted $2n$ -clique from which edges $C = \{(2i-1, 2i) \mid i = 1, \dots, n\}$ have been removed—with an example provided in Figure 4.3— \tilde{W}^{KC} is W^{KC} perturbed with k edges $S = \{(2i-1, 2i)\} \subseteq C$, $\mathbb{1} = \mathbb{1}_{2n}$, and \otimes is Kronecker product.

It is easy to show that the proposed input to DIVER is indeed legal: W^{KC} is row-stochastic matrix of a uniformly weighted undirected strongly connected aperiodic network; and $\tilde{x} = z \otimes \mathbb{1}_2$ and $x = \frac{s\mathbb{1} + m(z \otimes \mathbb{1}_2)}{m+k}$ are legal vectors of altered and original user opinions, respectively.

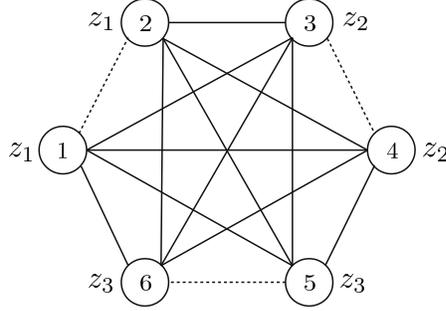


Figure 4.3: Network W^{KC} for $n = 3$; absent edges C are displayed dashed. Node states $z \otimes \mathbb{1}_2$, used in the reduction, are displayed next to the nodes.

Let us show what DIVER transforms into under the proposed input of (4.5). First, we notice that, for $b(k, x, \tilde{x}) = \langle \pi, (m+k)x - m\tilde{x} \rangle$ of (4.4), the following holds

$$b\left(k, \frac{s\mathbb{1} + m\tilde{x}}{m+k}, \tilde{x}\right) = \langle \pi, s\mathbb{1} + m\tilde{x} - m\tilde{x} \rangle = s\langle \pi, \mathbb{1} \rangle = s.$$

Then, DIVER's objective (4.4) under input (4.5) will look as

$$f'(\tilde{W}^{KC}) = \left| \left\langle \sum_{(i,j) \in S} (e_i + e_j), \frac{z \otimes \mathbb{1}_2}{2} \right\rangle - s \right| = \left| \sum_{\ell=1}^n y_\ell z_\ell - s \right|,$$

where y_ℓ are edge decision variables

$$y_\ell = \begin{cases} 1, & \text{if } (2\ell - 1, 2\ell) \in S, \\ 0, & \text{otherwise.} \end{cases}$$

Thus, solving DIVER via minimizing $f'(\tilde{W}^{KC})$, we look for a subset of $\{z_\ell\}$ of size k summing up to s , which is exactly what kSSP_{01} is after, so $\text{kSSP}_{01} \propto \text{DIVER}$.

Parts 1) and 4) of the proof together establish the reduction chain $\text{SSP} \propto \text{kSSP}_{01} \propto \text{DIVER}$, so DIVER is NP-hard. ■

DIVER's hardness is exacerbated by the fact that standard approximation techniques [173] are not applicable, as DIVER's objective function is not submodular.

4.4 Related Work

While DIVER is a new problem, there is a range of problems—in extremal network design as well as in the perturbation analysis of centrality measures and stationary distributions of Markov chains—related to ours either in the nature of the optimized objective or that of the methods and analyses involved. In this section, we survey several groups of these works.

4.4.1 Analytic Optimization of Network Topology

Here, we review network design problems, where a network’s topology is altered to optimize some analytic property of that network.

Algebraic Connectivity: Ghosh and Boyd [174] studied the problem of maximizing the algebraic connectivity—the second smallest eigenvalue λ_2 of the Laplacian—of an undirected unweighted network via edge addition. The authors use convex relaxation to formulate the problem as a semidefinite program (SDP), which is feasible to solve for small networks. They also provide a greedy perturbation heuristic that picks edges (i, j) based on the largest value of $(v_i - v_j)^2$ —the squared difference of the Fiedler vector’s components corresponding to each edge’s ends. The authors show that, in case of simple λ_2 , value $(v_i - v_j)^2$ gives the first-order approximation of the increase in λ_2 if edge (i, j) is added to the network. The authors also derive bounds on algebraic connectivity under single-edge perturbation. More recently, this approach has been employed by Yu et al. [175] for the design of an edge selection heuristic that the authors have augmented with an extra objective—neighborhood overlap-based user similarity (which likely correlates with edge acceptance likelihood).

Spectral Radius: Van Mieghem et al. [176] study the problem of minimizing the spectral radius of an undirected network via edge or node removal. They prove NP-hardness of the

problem, and show that the edge selection heuristic that picks edges (i, j) with the largest scores $v_i v_j$ —where v is the dominant eigenvector—performs well in practice. More recently, Saha et al. [177] addressed the same spectral radius minimization problem and designed a walk-based algorithm, relying on the link between the sum of powers of eigenvalues of a network and the number of closed walks in it, and provided approximation guarantees for it. Zhang et al. [178] studied spectral radius minimization for directed networks under SIR model, and provided an SDP/LP-based solution, having high polynomial time complexity.

Eigenvalues and Their Functions: Tong et al. [179] target optimization of the diffusion rate—expressed as the largest eigenvalue of the adjacency matrix—through a directed strongly connected unweighted (see [180] for the weighted case) network via edge addition or removal. Similarly to [176], the authors use first-order perturbation theory to assess the effect of deleting k edges $\lambda_{max} - \tilde{\lambda}_{max} = \sum u_i v_j / \langle u, v \rangle + O(k)$, where u and v are the left and right dominant eigenvectors of the adjacency matrix, respectively. This analysis inspires an edge selection heuristic, with the quality of edge (i, j) being defined as $u_i v_j$, similarly to $v_i v_j$ edge score of [176]. Le et al. [181] extend this result to the networks with small eigengaps. Chan et al. [182] target optimization of natural connectivity—a network robustness measure defined, roughly, as an average of exponentiated eigenvalues of the adjacency matrix—of an undirected network via altering its topology. For edge addition, they focus on the edges between high-centrality nodes.

Other Objectives: The SDP-based approach of Ghosh and Boyd [183] has been applied by the same authors to minimization of the total effective resistance of an undirected electric network via edge weight selection. Arrigo and Benzi [184] address the problem of optimizing the total communicability—the sum of the entries in the exponential of the adjacency matrix—in an undirected connected network via edge addition and removal. The authors use edge se-

lection heuristics, favoring edges between the nodes having high eigenvector centrality (for edge addition) or edges between the nodes having a large sum of their degrees (for edge removal). Garimella et al. [185] study an edge recommendation problem targeting reduction of polarization in a directed unweighted network, where polarization is measured via a random walk-based score. The edges are created between users “holding opposing views”. Similarly to [182, 184], the authors use an edge-selection heuristic that favors edges between high-degree nodes.

4.4.2 Combinatorial Optimization of Network Topology

These works address network design problems whose objectives or methods are of combinatorial nature. A large portion of these works are dedicated to direct information spread optimization in combinatorial opinion dynamics models, in contrast to indirectly optimizing some analytic feature of the network, such as the spectral radius of its adjacency matrix, expected to facilitate or hinder information propagation.

Information Spread: Chaoji et al. [186] look at a problem of maximizing the size of the activated user set under the Independent Cascade-like opinion dynamics model in an undirected network via edge addition. The authors prove NP-hardness of the problem, apply continuous relaxation to gain submodularity of the objective, and design a greedy cubic-time approximation algorithm for the relaxed problem. Kuhlman et al. [187] focus on general threshold-based propagation models, and address the problem of minimizing the contagion spread via edge deletion in a directed weighted network. The authors prove inapproximability of the problem, and design a spread simulation-based heuristic, that proves to be effective in experiments. The work of Khalil et al. [188] is dedicated to facilitating or hindering the spread of information under Linear Threshold (LT) model via edge addition or deletion in a directed weighted network. The authors design an influence objective function and prove its supermodularity. The

latter property used together with sampling of LT process realizations allows for the design of an efficient linear-time algorithm for target edge selection.

Shortest Paths and Optimal Network Flows: Phillips [189] studied the problem of minimizing a combinatorial maximum flow / minimum cut in a network. Each capacitated edge has a destruction cost, and the adversary needs to select a subset of edges to destroy, constrained by the total edge destruction budget. The authors prove NP-hardness of the problem, and design an FPTAS for the case of a planar network. Israeli and Wood [190] conduct a study of an NP-hard problem of maximizing a single s - t shortest path via edge removal in a directed network, formulated as a mixed-integer program (MIP). Due to the prohibitive time complexity of a direct solution of a MIP problem, the authors propose several decomposition techniques to accelerate the computation under some assumptions on the edge removal delays. Papageorgiou et al. [191] address the problem of minimizing the average all-pairs shortest path length in a connected undirected network via edge addition, and propose a greedy algorithm and two heuristics. Their most efficient algorithm has a quadratic time complexity. Ishakian et al. [192] define a general path-counting centrality measure and study a problem of maximizing the centrality of a given node via edge addition in a DAG. The authors use a quadratic-time greedy strategy for picking edges providing the largest marginal increase of the objective. Parotsidis et al. [193] study minimization of the sum of lengths of the shortest paths from a target node to all other nodes via link recommendation to the target node in an undirected network. The problem is proven to be NP-hard, and an efficient approximation algorithm is designed, employing submodularity of the objective. A related problem of minimizing the maximal shortest path length has been previously addressed by Perumal et al. [194]; another related problem of maximizing the coverage centrality—the number of unique node pairs whose shortest paths pass through a given node—is addressed by Medya et al. [195].

4.4.3 Centrality Perturbation and Manipulation

These works study either how eigenvector centrality, or PageRank, or the stationary distribution of a Markov chain changes when a network’s structure is perturbed, or how to strategically manipulate centrality by altering the network.

Strategic Centrality Manipulation: Avrachenkov and Litvak [196] analyze to what extent a node can improve its PageRank by creating new out-edges. The authors derive equalities that result in a conclusion that the PageRank of a web-page cannot be considerably improved by restructuring its out-neighborhood. The authors also derive an optimal linking strategy, stating that it is optimal for a web-page to have only one outgoing edge pointing to a web-page with the shortest mean first passage time back to the original page. Similar conclusions can be drawn for eigenvector centrality in an arbitrarily weighted network using Theorem 18 from Section 4.5.2 of this thesis. De Kerchove et al. [197] generalize the results of Avrachenkov and Litvak [196], studying maximization of the sum of PageRanks of a subset of nodes via adding outgoing edges to them. Csáji et al. [198] study the problem of optimizing the PageRank of a given node via directed edge addition. The authors formulate the optimization problem as a Markov decision process and propose a (generally, not scalable) polynomial-time algorithm for it.

Centrality Perturbation Analysis: Cho and Meyer [199] provide coarse bounds of type $|\pi_i - \tilde{\pi}_i|/\pi_i \leq \|E\|_\infty \max_{i \neq j} m_{ij}/2$ for the stationary distribution of a generally perturbed Markov chain, where E is an additive perturbation of the state transition matrix. Chien et al. [200] provide an efficient algorithm for incremental computation of PageRank over an evolving edge-perturbed graph, with the analysis’ drawing upon the theory of Markov chains. The key idea of their algorithm is to contract the network and localize its part where the nodes are likely to have changed their PageRank scores under the perturbation. Langville and Meyer [201] provide ex-

act equalities for the change in the stationary distribution of a perturbed Markov chain using group inverses. They address the problem of updating the stationary distribution under multi-row perturbation via exact and approximate aggregation, similarly to what Chien et al. [200] did for PageRank. Hunter [202] addresses the same problem of establishing equalities for the change in the stationary distribution, yet, provides an answer that does not involve group inverses and, instead, uses mean first passage times in a Markov chain; our perturbation analysis in Section 4.5.2 builds upon this result. Como and Fagnani [203] provide an upper bound on the perturbation of the stationary distribution of a Markov chain in terms of the mixing time of the chain as well as the entrance time to and the escape likelihood from the states with perturbed out-neighborhoods. Bahmani et al. [204] address the problem of updating PageRank algorithmically. The proposed node probing-based algorithms provide a close estimate of the network’s PageRank vector by crawling a small portion of the network. More recently, Li et al. [205] and Chen and Tong [206] addressed a general problem of updating eigenpairs of an evolving network. Chen and Tong provide a linear-time algorithm for tracking top eigenpairs.

4.5 Strategic Edge Addition to the Network

Since, due to Theorem 17, DIVER optimization problem

$$\text{DIVER}(W, k, x, \tilde{x}) = \arg \min_{\tilde{W}} |\langle \tilde{\pi}(\tilde{W}), \tilde{x} \rangle - \langle \pi, x \rangle|, \quad (4.1)$$

is NP-hard, we need to design a heuristic for it. Our general approach—formalized later in Section 4.5.5—is as follows. We will assess candidate edges with respect to how much their addition to the network can decrease term $\langle \tilde{\pi}, \tilde{x} \rangle$ of (4.1), and then iteratively add the most promising edges to the network until we are satisfied with the value of DIVER’s objective.

Thus, our foremost concerns now are: the selection of a small number of candidate edges

to assess, and the subsequent assessment of the potential impact of these edges’ addition to the network upon the network’s eigenvector centrality. These concerns are addressed in the following two sections.

4.5.1 Selection of Candidate Edge Source Nodes

The above described general approach involves assessing candidate edges individually. However, the number of absent edges in a sparse network is $O(n^2)$, and inspecting all of them is unfeasible for large networks. Hence, we will focus on a small number of candidate edges, outgoing from $n_{src} \ll n$ network nodes. This implies that a small number of nodes are being the sources for most—or, at least, a large number of—“good” candidate edges. Intuitively, the nodes having the largest (eigenvector) centrality should be those edge sources; the changes in their out-neighborhoods should have the largest impact upon the centrality distribution in the network, as Figure 4.4 suggests. This intuition will find formal support in Corollary 3 in the following Section 4.5.2.

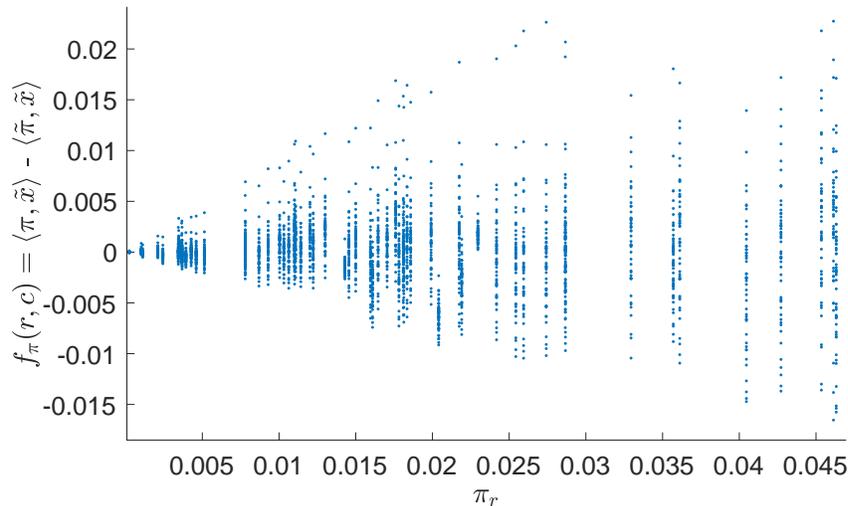


Figure 4.4: Dependence of the asymptotic consensus value’s reduction $f_{\pi}(r, c) = \langle \pi, \tilde{x} \rangle - \langle \tilde{\pi}, \tilde{x} \rangle$ after addition of edge (r, c) , $\theta_{rc} = const$ to a scale-free network ($n = 100, \gamma = -2.5$) upon the eigenvector centrality π_r of the edge’s source node.

Thus, to make sure that addition of candidate edges to the network has a large impact—either very positive or very negative—upon the asymptotic consensus value, we can limit ourselves to considering only candidate edges outgoing from high-centrality nodes. Fortunately, the number of such nodes in real-world social networks is indeed small, and most nodes are at the periphery, which justifies our choice of $n_{src} \ll n$.

4.5.2 Eigencentality Under Single-Edge Perturbation

In order to tackle DIVER (4.1), we need to understand how the addition of a single edge $\tilde{w}_{rc} = \theta_{rc} \in (0, 1]$ from node r to node c changes the network's eigenvector centrality $\pi \rightarrow \tilde{\pi}$. We assume that edge (r, c) is originally absent— $w_{rc} = 0$ —and use the same single-edge perturbation model

$$\tilde{W} = W - \theta_{rc} \text{diag}(e_r)W + \theta_{rc} e_r e_c^T. \quad (4.2)$$

In our subsequent perturbation analysis, we will make the following Assumption 4.

Assumption 4 (Rational Selfishness). *Users are rationally selfish in that for any user i , $\forall j \neq i$: $w_{ii} > w_{ij}$. Thus, each user trusts his or her own opinion more than the opinion of any other individual user.*

The following Theorem 18 states how the eigencentality vector changes under a single-edge perturbation (4.2).

Theorem 18 (Single-Edge Perturbation). *Under Assumption 4, for a single-edge perturbation (4.2) of a strongly connected aperiodic network with adjacency matrix W , the network's eigenvector centrality changes as*

$$\tilde{\pi}_j = \pi_j \left[1 - \frac{\theta_{rc}(m_{cj}(1 - \delta\{j, c\}) - m_{rj} + 1)}{m_{rr} + \theta_{rc}(m_{cr} - m_{rr} + 1)} \right], \quad (4.6)$$

where m_{ij} is the mean first passage time from state i to j of Markov chain W , and δ is Kronecker delta. In particular,

$$\tilde{\pi}_r = 1/[m_{rr} + \theta_{rc}(m_{cr} - m_{rr} + 1)]. \quad (4.7)$$

The proof will rely on the following Theorem 19 due to Hunter [202], provided for reference below.

Theorem 19 ([202, Theorem 4.4]). *Let multiple perturbations occur in r 'th row of W . Let $\varepsilon_i = \tilde{W}_{ri} - W_{ri}$; additionally, let the minimal negative perturbation happen at state a , with $\varepsilon_a = -m = \min \{\varepsilon_j \mid 1 \leq j \leq n\}$; and the maximal positive perturbation occur at state b with $\varepsilon_b = M = \max \{\varepsilon_j \mid 1 \leq j \leq n\}$. Also, let P be the set of positive perturbation indices, excluding b , and N be the set of negative perturbation indices, excluding a . Then,*

$$\pi_j - \tilde{\pi}_j = \begin{cases} \pi_a \tilde{\pi}_r [Mm_{ba} + \sum_{k \in P \cup N} \varepsilon_k m_{ka}] & \text{if } j = a, \\ \pi_b \tilde{\pi}_r [-mm_{ab} + \sum_{k \in P \cup N} \varepsilon_k m_{kb}] & \text{if } j = b, \\ \pi_j \tilde{\pi}_r [-mm_{aj} + Mm_{bj} + \sum_{\substack{k \in P \cup N \\ k \neq j}} \varepsilon_k m_{kj}] & \text{if } j \neq a, b. \end{cases}$$

Proof (Theorem 18). Let us apply Theorem 19 to our case of a single-edge perturbation (4.2). We are adding edge (r, c) with weight θ_{rc} to the network. Due to the form (4.2) of our single-edge perturbation, the only positive perturbation occurs at the added edge's destination node c , so $b = c$, $\varepsilon_c = M = \theta_{rc}$, and $P = \emptyset$. For all the other out-neighbors i of the new edge's source node r , the corresponding perturbations $\varepsilon_i = -\theta_{rc}w_{ri}$ are negative. Due to Assumption 4, $\forall i \neq r : w_{rr} > w_{ri}$, so the minimal negative perturbation occurs at $i = r$, and, thus, $a = r$ and $\varepsilon_a = -m = -\theta_{rc}w_{rr}$.

Let us first show the validity of (4.6) in case of $j = r$, that is, (4.7). According to Theorem 19,

$$\begin{aligned}\pi_r - \tilde{\pi}_r &= \pi_r \tilde{\pi}_r \left[\theta_{rc} m_{cr} + \sum_{\substack{k \in \text{PUN} \\ k \neq r}} (-\theta_{rc} w_{rk}) m_{kr} \right] \\ &= (\text{as } w_{rc} = 0) = \theta_{rc} \pi_r \tilde{\pi}_r \left[m_{cr} - \sum_{k \neq r} w_{rk} m_{kr} \right].\end{aligned}$$

Using the one-hop conditioning Theorem 16, the obtained expression can be written as

$$\begin{aligned}\pi_r - \tilde{\pi}_r &= \theta_{rc} \pi_r \tilde{\pi}_r [m_{cr} - m_{rr} + 1] \\ \Leftrightarrow \tilde{\pi}_r &= \pi_r / (1 + \theta_{rc} \pi_r [m_{cr} - m_{rr} + 1]).\end{aligned}$$

Dividing the numerator and denominator in the right-hand side of the obtained expression by $\pi_r > 0$ —where positivity comes from Perron-Frobenius theorem—and using equality $1/\pi_r = m_{rr}$ (from Theorem 15), we obtain (4.7).

Let us similarly deal with the case $j \neq r, c$. From Theorem 19,

$$\begin{aligned}\pi_j - \tilde{\pi}_j &= \pi_j \tilde{\pi}_r \left[-\theta_{rc} w_{rr} m_{rj} + \theta_{rc} m_{cj} + \sum_{\substack{k \in \text{PUN} \\ k \neq j}} (-\theta_{rc} w_{rk}) m_{kj} \right] \\ &= \theta_{rc} \pi_j \tilde{\pi}_r \left[m_{cj} - w_{rr} m_{rj} - \sum_{k \neq r, c, j} w_{rk} m_{kj} \right] \\ &= (\text{as } w_{rc} = 0) = \theta_{rc} \pi_j \tilde{\pi}_r \left[m_{cj} - \sum_{k \neq j} w_{rk} m_{kj} \right] \\ \Leftrightarrow (\text{from Theorem 16}) &\Leftrightarrow \tilde{\pi}_j = \pi_j [1 - \theta_{rc} \tilde{\pi}_r (m_{cj} - m_{rj} + 1)].\end{aligned}$$

Substituting (4.7) in the obtained expression, we get (4.6) for $j \neq c$. The proof for $j = c$ is similar and, thus, is omitted. ■

The following Corollary 3—justifying Section 4.4’s focus on top-centrality edge source nodes—immediately follows from equation (4.6) of Theorem 18 used with Theorem 15.

Corollary 3. *Under perturbation (4.2) of the network with a single edge (r, c) , $\theta_{rc} > 0$, it holds that $\lim_{\pi_r \rightarrow 0} \tilde{\pi} = \pi$, and, thus, $\lim_{\pi_r \rightarrow 0} f_\pi(r, c) = \lim_{\pi_r \rightarrow 0} (\langle \pi, \tilde{x} \rangle - \langle \tilde{\pi}, \tilde{x} \rangle) = 0$.*

4.5.3 Asymptotic Consensus Value Under Single-Edge Perturbation

To solve DIVER, we are interested in adding candidate edges that would result in a large reduction $f_\pi(r, c) = \langle \pi, \tilde{x} \rangle - \langle \tilde{\pi}, \tilde{x} \rangle$ of the asymptotic consensus value. While Theorem 18 states how different components of the eigencentality vector change under a single-edge perturbation (4.2), the following Theorem 20 is concerned with the effect of such perturbation upon the value of $f_\pi(r, c)$. The proof of the theorem is immediately obtained by substituting (4.6) of Theorem 18 into $f_\pi(r, c) = \langle \pi - \tilde{\pi}, \tilde{x} \rangle$.

Theorem 20. *Under the rational selfishness Assumption 4, for a single-edge perturbation (4.2) of W , the asymptotic consensus value $\langle \pi, \tilde{x} \rangle$ decreases as follows:*

$$f_\pi(r, c) = \langle \pi, \tilde{x} \rangle - \langle \tilde{\pi}, \tilde{x} \rangle = \theta_{rc} \frac{\sum_{j=1}^n \pi_j (m_{cj} \cdot (1 - \delta\{j, c\}) - m_{rj} + 1) \tilde{x}_j}{m_{rr} + \theta_{rc} (m_{rc} - m_{rr} + 1)}. \quad (4.8)$$

Essentially, Theorem 20 provides us with an edge score $f_\pi(r, c)$, whose use for candidate edge selection comprises our heuristic for DIVER. Unfortunately, f_π ’s computation is rather challenging, and is addressed in the following section.

4.5.4 Efficient Computation of Candidate Edge Scores

Computation of candidate edge scores f_π is challenging for two reasons. Firstly, expression (4.8) involves summation over all n network nodes. Since there are $O(n_{src}n)$ candidate edges—with $n_{src} \ll n$ sources and n destinations—it would result in at least a quadratic-time heuristic for DIVER that would not scale. Secondly, expression (4.8) involves mean first passage times, whose direct computation is very expensive. We address both these challenges separately below.

Focus on a Small Number of Nodes: Our first concern is that expression (4.8) for f_π contains summation over all n nodes. Intuitively, not all network nodes contribute equally to the value of (4.8). Indeed, in networks with skewed eigencentality distribution, such as scale-free networks, f_π is largely determined by the top-centrality nodes. This is illustrated in Figure 4.5, that shows the relationship between the exactly computed f_π and its approximations, with different numbers of scale-free network nodes being used in f_π 's computation. We can see that,

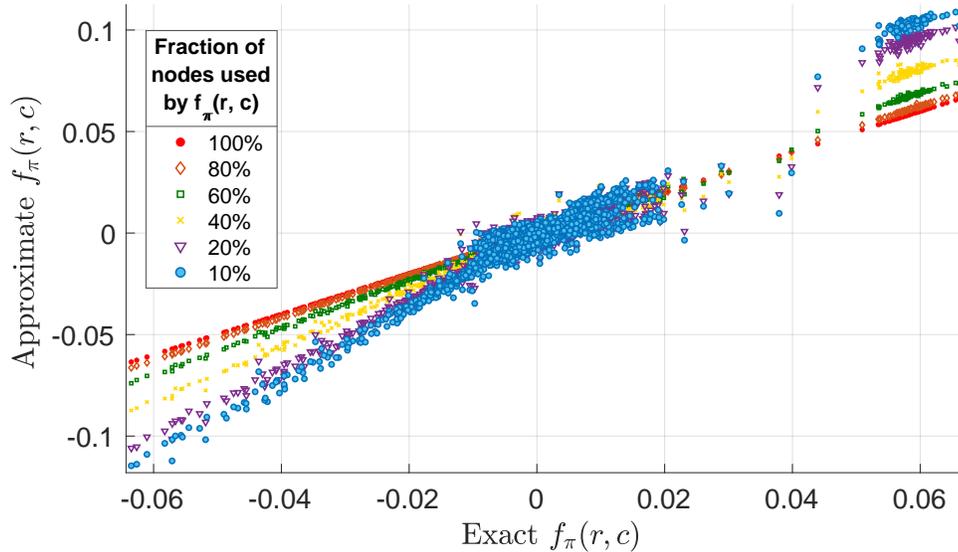


Figure 4.5: Comparison of exact and approximate candidate edge scores f_π in a scale-free network ($n = 100, \gamma = -2.5$).

even when we use only 10% of nodes, the relative order of f_π for different candidate edges is

close to the original, and it is still easy to identify candidate edges (r, c) with the largest values of $f_\pi(r, c)$.

Thus, to efficiently compute $f_\pi(r, c)$, we will use only those j in (4.8) corresponding to a constant number, e.g. n_{src} , of top-centrality nodes in the network, in addition to $j \in \{r, c\}$.

Efficient Computation of Mean First Passage Times: In the previous section, we have considerably simplified computation of $f_\pi(r, c)$ by leaving only $O(n_{src})$ summands in expression (4.8). Now, our concern is to actually find values of the mean first passage times m_{ij} remaining in (4.8).

The classic method for exact MFPT computation [168, Theorem 4.4.7] involves computing the fundamental matrix $Z = (I - W + \mathbb{1}\pi^\top)^{-1}$ of Markov chain W , and defines MFPTs as

$$M = \{m_{ij}\} = (I - Z + \mathbb{1}\mathbb{1}^\top \text{diag}(Z)) \text{diag}^{-1}(\pi).$$

Computation of the fundamental matrix involves a cubic-time matrix inversion and would not scale. In [207], Hunter surveys alternative methods for MFPT computation, but all of them share the same high complexity. Most importantly, however, all existing methods target computation of all $O(n^2)$ MFPTs between all the nodes in the network.

Let us notice that expression (4.8) for f_π uses MFPTs either from or to high-centrality nodes: r are top-centrality according to Section 4.5.1; j are top-centrality according to Section 4.5.4. There are $n_{src}n \ll n^2$ such MFPTs, where n_{src} is the number of candidate edge source nodes r .

We propose to estimate the MFPTs between a small number of nodes by performing a finite random walk over the network and tracking first passage times between the nodes. The walk starts at an arbitrary node and proceeds for a predefined number of steps following the transition probabilities defined by the adjacency matrix W , viewed here as the state transition matrix of

a Markov chain. While performing the walk, we accumulate the passage times between n_{src} candidate edge sources and n candidate edge destinations, and compute the means when the walk is complete. This approach towards MFPT estimation is similar to the k -Step Markov Approach that White and Smyth [208, Section 6.4] used for estimation of their MFPT-based relative importance of network nodes.

The key questions here are *whether the proposed method will produce good estimates of MFPTs to/from high-centrality nodes* and, *if so, how long the random walk should be*. We answer these two questions via empirical analysis.

The first insight is that MFPTs to and from high-centrality nodes converge very fast, since the walk visits such nodes most often. This is illustrated in Figure 4.6, according to which the error of MFPT estimates m_{ij} noticeably varies with the walk’s length when both i and j are low-centrality, and is uniformly low if at least one of i and j is high-centrality.

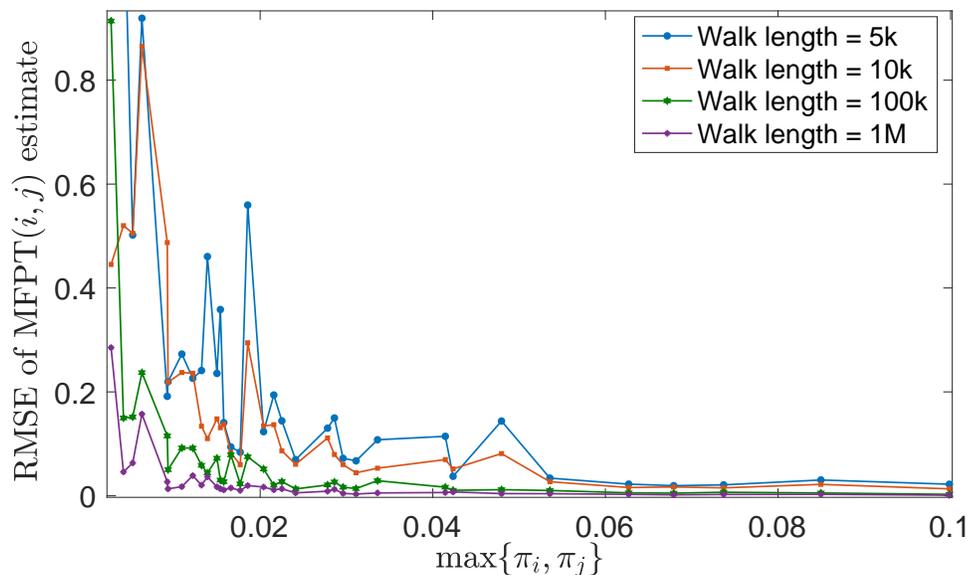


Figure 4.6: Root mean square error of MFPT estimates in a scale-free network ($n = 100, \gamma = -2.5$) using walks of different length.

The latter insight echoes the result of Avrachenkov et al. [209], who show that PageRank of high-centrality nodes estimated via Monte Carlo simulation converge very fast.

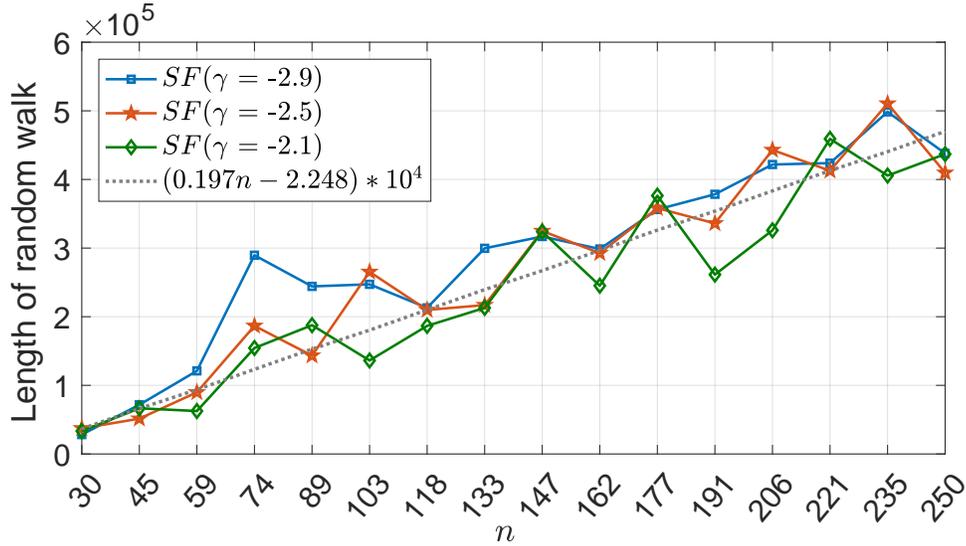


Figure 4.7: Dependency of the length of a random walk—used for estimating MFPTs to and from top-centrality nodes—on the size and density of the scale-free network.

Now, we empirically study the question of how long the random walk should be to obtain sufficiently good estimates of MFPTs to and from top-centrality nodes in a scale-free network. The results are reported in Figure 4.7, that shows how many steps a random walk should perform in order for 5% of MFPTs to and from top 5% high-centrality nodes to converge within 5% of their true values, while the network’s size n and scale-free exponent γ vary. For each pair (n, γ) , 100 networks are generated, and the mean walk lengths are reported. The results are reported for 3 specific scale-free exponents, $\gamma \in \{-2.9, -2.5, -2.1\}$. The length of the random walk does not depend on the scale-free exponent, and depends upon the network’s size n as $(0.197n - 2.248) \cdot 10^4$. This result allows to make the following statement.

Proposition 1 (Random Walk Length). *In scale-free networks, the length of a random walk sufficient for convergence of $O(n)$ MFPTs to and from $O(1)$ top-centrality nodes is $O(n)$ (in contrast to $O(n^3)$ cost of the direct computation of all MFPTs via the fundamental matrix method).*

4.5.5 Solving DIVER

In this section, we gather all our results, formally state a heuristic for solving DIVER, and analyze its complexity.

Algorithm 1 Heuristic for DIVER

Input: W —row-stochastic irreducible aperiodic sparse interpersonal appraisal matrix; k —number of new edges to add; n_{src} —maximal number of new edges’ sources.

Output: sequence $(r_1, c_1), (r_2, c_2), \dots$ of new edges to add

- 1: Compute eigenvector centrality π
 - 2: Define candidate edge source nodes:
 $R \leftarrow n_{src}$ top-centrality nodes in the network
 - 3: Estimate MFPTs $\{m_{ij}\}$ to and from each $r \in R$
 - 4: **for** $r \in R, c \in \{1, \dots, n\}$ **do**
 - 5: Estimate $f_\pi(r, c)$ using $O(n_{src})$ top-centrality nodes
 - 6: **end for**
 - 7: $S \leftarrow$ candidate edges (r, c) having top- k scores $f_\pi(r, c)$
 - 8: **return** S
-

Theorem 21. *Time-complexity of Algorithm 1 is $O(n(\text{gap}(W) + n_{src}^2) + n_{src} \log n_{src} + k \log k)$, where $\text{gap}(W)$ is the number of matrix-vector multiplications the power method uses to compute the dominant left eigenvector of W .*

Proof. In step 1 of Algorithm 1, we compute the dominant left eigenvector π of W using the power method, performing $\text{gap}(W)$ matrix-vector multiplications, each of which has a linear time complexity for sparse W . Thus, this step’s complexity is $T_1 = O(\text{gap}(W)n)$. The cost of selecting top n_{src} elements out of n at Step 2 is $T_2 = O(n + n_{src} \log(n_{src}))$. In step 3, following Section 5.5.2 and, in particular, Proposition 5.1, we estimate MFPTs via a $O(n)$ -long finite random walk, so this step’s cost is $T_3 = O(n)$. At steps 4-6, we compute $n_{src}n$ edge scores f_π . Following the method of Section 5.5.1, each $f_\pi(r, c)$ is computed in time $O(n_{src})$, bringing time complexity of steps 4-6 to $T_{4-6} = O(n_{src}^2 n)$. Finally, selection of top k out of $n_{src}n$ items at step 7 is performed in time $T_7 = O(n_{src}n + k \log k)$. If we collect the expressions for T_1, \dots, T_7 , we get $T = O(n(\text{gap}(W) + n_{src}^2) + n_{src} \log n_{src} + k \log k)$. ■

In Theorem 21, the number $\text{gap}(W)$ of iterations it takes the power method to converge depends on W 's spectral gap λ_2/λ_1 , but, in practice, $\text{gap}(W)$ usually can be assumed to be a reasonably small constant [209]. Thus, assuming that $\text{gap}(W)$ is bounded, as well as noticing that we choose both n_{src} and k to be small, that is, $n_{src} \ll n$ and $k \ll n$, it immediately follows from Theorem 21 that *Algorithm 1 is computable in time $O(n)$* .

4.5.6 Probabilistic Edge Addition

Our theory and algorithms described so far straightforwardly extend to the case of probabilistic edge recommendation. Let us assume candidate edge (r, c) is accepted with probability $\mathbb{P}_{rc} \in [0, 1]^1$. Then, due to linearity of the inner product operation, the probabilistic version $f_{\pi}^{\mathbb{P}}(r, c)$ of the edge score $f_{\pi}(r, c) = \langle \pi - \tilde{\pi}, \tilde{x} \rangle$ is

$$f_{\pi}^{\mathbb{P}}(r, c) = \mathbb{P}_{rc} \langle \pi - \tilde{\pi}, \tilde{x} \rangle + (1 - \mathbb{P}_{rc}) \langle \pi - \pi, \tilde{x} \rangle = \mathbb{P}_{rc} f_{\pi}(r, c),$$

and Algorithm 1 (lines 5, 7) changes in that it computes $f_{\pi}^{\mathbb{P}}(r, c)$ instead of $f_{\pi}(r, c)$ using the same procedure from Section 4.5.4. We can further change Algorithm 1 (line 2) to select edge sources not just based on their eigencentality π_i , but rather based on the value of the product $\pi_i \max_j \mathbb{P}_{ij}$, thereby, avoiding recommending new edges to “celebrity nodes” having $\max_j \mathbb{P}_{ij} \approx 0$. Notice that the above proposed alternations do not affect Algorithm 1’s time complexity.

4.6 Experimental Results

In this section, we experimentally study Algorithm 1’s performance on synthetic and real-world networks. We start with an experimental setup, and, then, study the performance of our

¹ \mathbb{P}_{rc} can be estimated based on the present links and the history of their creation. Alternatively, we can choose $\mathbb{P}_{rc} \in \{0, 1\}$ ourselves, solving a problem where only a subset of candidate edges is available for recommendation.

heuristic, as well as its effectiveness, efficiency, and robustness to network noise.

4.6.1 Experimental Setup

Networks: We have experimented with three synthetic and three real-world networks:

- **SF**(n, γ): a scale-free network with n nodes and scale-free exponent γ , built using a greedy generator that constructs a directed graph trying to match a given degree distribution².
- **BA**(n, m_{link}): a Barabási-Albert network [45] on n nodes with seed SF($0.01n, -2.5$) and m_{link} edges created per node / iteration.
- **ER**(n, \mathbb{P}_{edge}): an Erdős–Rényi network [44] with the edge probability \mathbb{P}_{edge} , being an example of a “non-scale-free-like” network, on which our heuristic and baselines perform poorly.
- **Karate**: Zachary’s Karate Club network [164].
- **Facebook**: a 4k-node part of Facebook graph [210].
- **Epinions**: a 32k-node part of the mutual trust network epinions.com [211].

The characteristics of these networks are summarized in Table 4.2, where S_π and K_π are skewness and kurtosis, respectively, of the unweighted network’s eigencentality distribution, and S_d and K_d are the same metrics of the same network’s total (in- plus out-) degree distribution.

All the networks except Facebook are directed. If the original network was not strongly connected, we replace it with its largest strongly connected component. We also add all self-loops—since each user in our case is supposed to have some amount of trust in her or his

²EvaluateGraphCreateRandomGraph.cpp of Complex Networks Package.

Network	$ V $	$ E $	S_π	S_d	K_π	K_d
SF(1024, -2.5)	1024	3.9k	5.30	9.37	54.88	134.56
BA(1024, 3)	1024	7.9k	7.51	7.46	103.34	84.26
ER(1024, 0.25)	1024	261k	0.04	0.10	2.95	3.01
Karate	34	190	1.09	2.00	3.24	6.30
Facebook	4039	181k	4.24	4.52	19.67	57.56
Epinions	32k	476k	9.87	9.02	130.01	160.64

Table 4.2: Networks used in experiments with DIVER.

own opinion—and draw edge weights uniformly at random, while maintaining self-weight dominance—as per Assumption 4—and rescaling the weights in each out-neighborhood to make sure the resulting adjacency matrix is row-stochastic.

Opinions and Attack Upon Them: The opinions x before the attack are drawn uniformly at random from $[0, 1]^n$. Then, to create \tilde{x} , $k_{bribed} = 0.1n$ (5 for Karate network) nodes are selected uniformly at random, and their opinions change to 1, while the opinions of other users stay intact.

Methods: We experiment with several versions of our heuristic, as well as two baselines.

- DIVER: uses Algorithm 1 with f_π being computed using *all* rather than $O(1)$ top-centrality nodes, and estimates MFPTs using the procedure of Section 4.5.4. The heuristic adds k edges at a time, improving the distribution of eigencentralities $\tilde{\pi}$ either until the asymptotic consensus value $\langle \tilde{\pi}, \tilde{x} \rangle$ gets close enough to its original state $\langle \pi, x \rangle$, or a maximum of k_{max} new edges is reached. The number n_{src} of top-centrality nodes considered in computation is 20 for synthetic, 10 for Karate, 40 for Facebook, and 200 for Epinions networks. This version of Algorithm 1 has quadratic time complexity, and represents “the best case” behavior of DIVER heuristic.
- DIVER($f_\pi \sim X\%$), $X \in \{10, 20, 30\}$: similar to DIVER above, except that Algorithm 1 estimates f_π using only a fraction X of all the nodes. The time complexity of these

methods is pseudo-linear, as per Theorem 21, and lower values of X correspond to a higher computational efficiency, yet, to a lower effectiveness (more new edges are spent to recover the original asymptotic consensus value).

- $\text{BASE}(\text{rnd})$: the worst-case baseline, that selects new edges uniformly at random.
- $\text{BASE}(\theta_{rc}(\pi_r - \pi_c)(\tilde{x}_r - \tilde{x}_c))$: this baseline attempts to effectively reduce the asymptotic consensus value $\langle \pi, \tilde{x} \rangle$ by adding heavy-weight edges from higher-centrality nodes r having larger opinion values to lower-centrality nodes c having lower opinion values, thereby, reducing the contribution of $\pi_r \tilde{x}_r$ and increasing that of $\pi_c \tilde{x}_c$ to the value of $\langle \pi, \tilde{x} \rangle$. This baseline—similarly to DIVER—ranks only the candidate edges outgoing from n_{src} top-centrality nodes, and, hence, is computable in pseudo-linear time (“pseudo” since it uses π). Conceptually, this baseline is “weaker” than DIVER in that it relies only on absolute node importance, while DIVER exploits relative node importance to assess candidate edges.

Evaluation: For all these methods, we assess their performance based on how many candidate edges a method uses to recover the asymptotic consensus value.

4.6.2 Solving DIVER

In this section, we solve DIVER using the previously described networks and methods. All these methods add $k = 10$ edges at a time, improving the distribution of eigencentralities $\tilde{\pi}$ either until the asymptotic consensus value $\langle \tilde{\pi}, \tilde{x} \rangle$ gets close enough to its original state $\langle \pi, x \rangle$, or a maximum of $k_{max} = 200$ new edges is reached. The results are displayed in Figure 4.8.

We can see that our heuristic works well on all scale-free-like networks, disabling attacks by adding 20-60 candidate edges. In case of ER network, the quality of all absent edges is uniformly low, which naturally leads to all the methods’ performing uniformly poorly. For

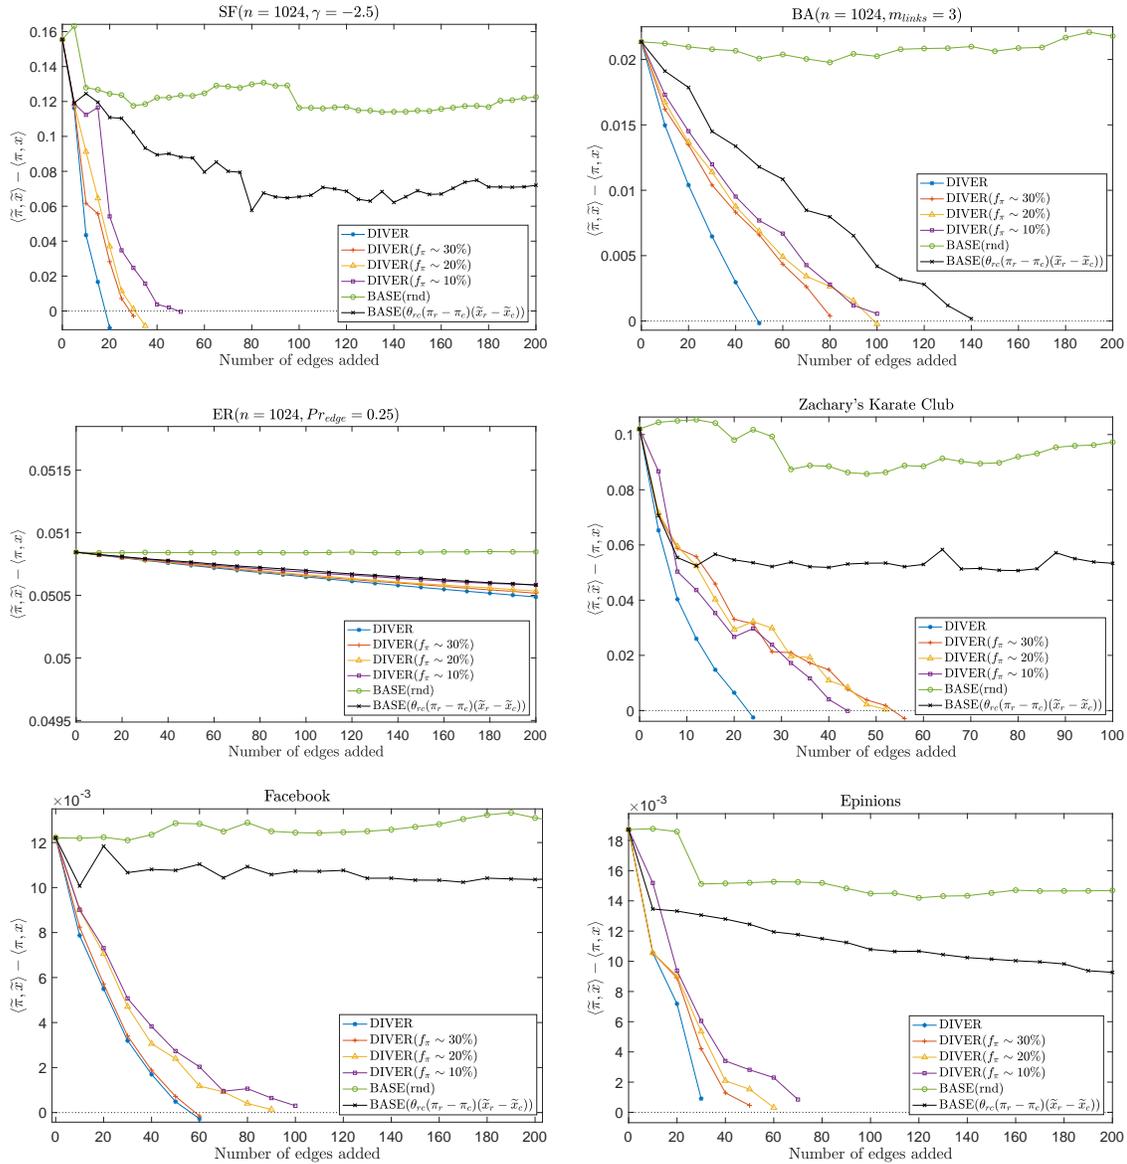


Figure 4.8: Solving DIVER over synthetic and real-world networks.

the case of Karate network—which, similarly to ER network, has rather low skewness and kurtosis of eigencentality and degree distributions as per Table 4.2—performance of DIVER heuristic degrades rather fast with the number of nodes used in computation of f_π ; for example, using 30% of nodes to compute f_π , it takes DIVER($f_\pi \sim 30\%$) twice as many new edges as DIVER to restore the asymptotic consensus value. We also notice that baseline BASE($\theta_{rc}(\pi_r - \pi_c)(\bar{x}_r - \bar{x}_c)$)

$\pi_c)(\tilde{x}_r - \tilde{x}_c))$ performs well only on BA network, and underperforms on all the other networks. The latter is expected, as, unlike DIVER, this baseline uses only the information about absolute node importance for candidate edge assessment.

4.6.3 DIVER’s Effectiveness and Efficiency

In the previous section, we observed that performance of $\text{DIVER}(f_\pi \sim X\%)$ deteriorates when X gets small. Here, we illustrate the impact of lower X upon the heuristic’s effectiveness.

Figure 4.9 shows the quality f_π of each candidate edge of $\text{SF}(256, -2.5)$ with $n_{src} = 20$, as well as the quality of edges selected by $\text{DIVER}(f_\pi \sim 20\%)$ over multiple iterations. While due

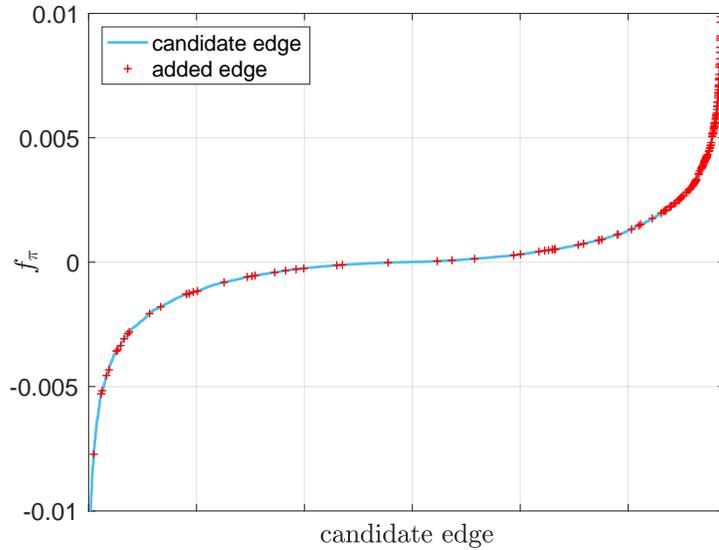


Figure 4.9: Quality of candidate edges added by $\text{DIVER}(f_\pi \sim 20\%)$ to $\text{SF}(256, -2.5)$.

to the approximate computation of f_π , some poor-quality edges are selected—a few of which actually drive the asymptotic consensus value up—most of the selected candidate edges are high-quality.

Figure 4.10 shows how the heuristic’s performance improves when we throw in more nodes into the computation of f_π .

Figure 4.11 shows the actual time it takes to perform a single evaluation of $\text{DIVER}(f_\pi \sim$

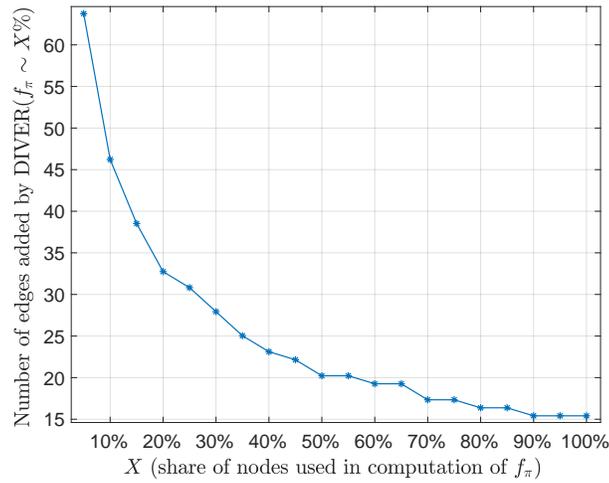


Figure 4.10: Dependence of $\text{DIVER}(f_\pi \sim X\%)$ performance on X over $SF(256, -2.5)$.

$100\% \frac{n_{src}}{n}$) adding $k = 10$ edges to $SF(1024, -2.5)$ using an Intel i7-5930K CPU. The heuristic

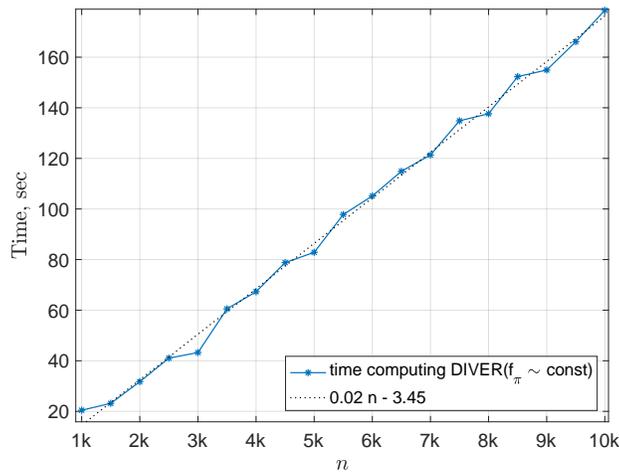


Figure 4.11: Time for a single evaluation of our heuristic $\text{DIVER}(f_\pi \sim const)$, adding $k = 10$ edges.

scales linearly, as per Theorem 21. However, if we extrapolate, a single evaluation of DIVER on a 100M-node network would take around 18 days. Thus, to make the method practical for such massive networks, we can improve its performance even further by considering only some of n nodes as destinations for the candidate edges; for example, we can consider only the destination nodes 2 hops away from the source nodes, increasing the acceptance likelihood of the recommended edges.

4.6.4 DIVER’s Robustness to Network Noise

We have studied what happens to performance of $\text{DIVER}(f_\pi \sim 100\%)$ when we do not know the exact weights of the edges in the network. We have run DIVER over $SF(1024, -2.5)$ network when edge weights are known exactly as well as when they deviate from their true values by up to 1% and up to 10%, respectively. The results are reported in Figure 4.12.

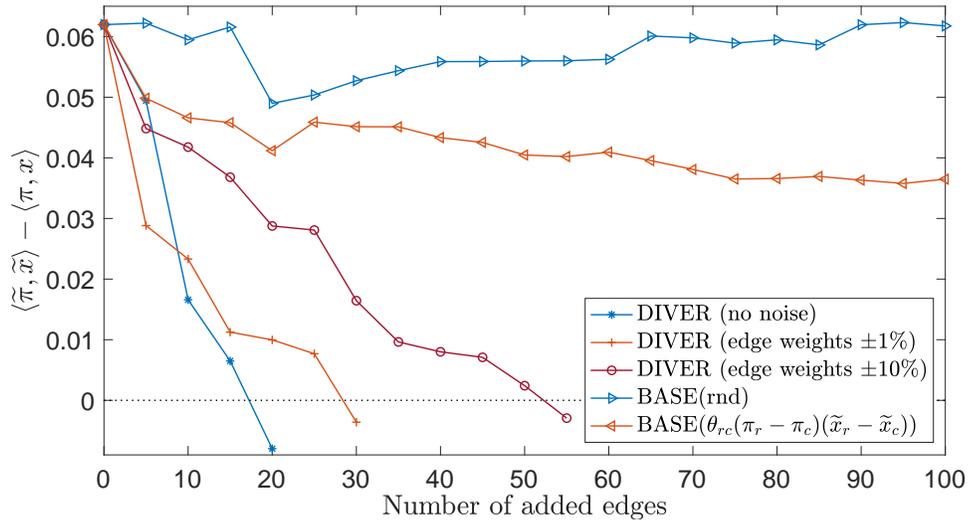


Figure 4.12: Performance of $\text{DIVER}(f_\pi \sim 100\%)$ over $SF(1024, -2.5)$ with exact and noisy edge weights.

We observe that DIVER’s performance expectedly worsens when we perturb the network, but still remains acceptable.

4.7 Conclusion

In this chapter, we defined DIVER—a new problem of strategically recommending edges in a social network to disable the effect of malicious external influence upon user opinions. Having shown that this problem is NP-hard, we focused on designing a heuristic for it. To that end, relying on the theory of Markov chains, we have provided a perturbation analysis, formally answering the question of how the network nodes’ eigencentralities and, thus, DIVER’s

objective change when an edge is added to the network. This analysis led to the definition of candidate edge scores that quantify the potential impact of candidate edges, allowing to add them to the network in a greedy fashion. We also provided insights into how to compute these edge scores in scale-free-like networks in pseudo-constant time, which resulted in a pseudo-linear-time heuristic for DIVER. One of these insights is related to efficient estimation of mean first passage times in Markov chains. We confirmed our theoretical and algorithmic findings in experiments with synthetic and real-world networks. Our results are rather general, and can be applied to other problems of strategic eigencentality manipulation.

4.8 Future Work

This chapter opens many avenues for potential *future research*, including the following.

- DIVER targets optimization of the scalar asymptotic consensus value, being the sum of the user opinions weighted by the corresponding users' eigenvector centralities. It may be fruitful to generalize DIVER to the optimization of the whole asymptotic opinion distribution, rather than its scalar summary.
- For the purposes of efficiently computing potential impact of candidate edge addition, we have studied the question of how to efficiently estimate mean first passage times to/from top-centrality states in a Markov chain. While our answer to the latter question was based on an empirical study on scale-free networks, providing formal convergence bounds for MFPT estimation would benefit many scientific areas dealing with Markov processes.
- Finally, while DIVER was proposed as a method to defend a social network against social control, it clearly can be used as an influence tool. Thus, it is important to study the ways to identify whether an edge recommendation process in a social network targets strategic change of the opinion distribution, as well as identify the goal of that process.

Chapter 5

Conclusion

In this thesis, we outline the scope of the emerging cross-discipline of dynamic processes in networks, and make contributions to the development of that field along the three dimensions of analysis, modeling, and control. In terms of analysis, relying upon combinatorial network algorithms, we devise a general method to analyze an observed process of polar opinion formation in social networks, with applications to anomaly detection in and extrapolation of a series of the process' states. With respect to modeling, we design a class of non-linear models capturing complex polar opinion formation behavior in social networks as well as connected to such physical processes as non-linear heat diffusion, and provide a general approach towards theoretical analysis of such models' behavior relying on the dynamical systems theory and non-smooth analysis. Finally, we address the problem of fighting social control via strategically augmenting social networks, building upon the theory of Markov chains, with our theoretical and algorithmic results' generally contributing to the area of the networked control of Markov processes.

This thesis opens up a number of avenues for future research, particularly, in dealing with processes in socio-economic networks.

5.1 Open Problems in Network Process Analysis

Below, we provide two general problems dealing with the analysis of observed network processes.

Predicting Future Evolution of Network Process: One open problem within the scope of the general network process analysis framework is that of the *efficient network process state series extrapolation*. Consider a situation when, having observed a network process' dynamics, we need to predict how it will evolve in the near future. For example, having observed how the political opinions of a social network's users have evolved in the past, we want to predict how these opinions will change at the time of the upcoming elections. In other words, we want to predict the future state of a network process based on a series of its past states. While the method developed in Chapter 2 of this thesis can estimate the likelihood of a specific change in a network process' state, here we are faced with the need to *search for the most plausible future state of the process*. One approach is to extrapolate the series of distances between the adjacent past states of the process to estimate the expected distance to the process' future state, and, then, among all future state candidates take as the prediction the state the distance to which is closest to the made distance estimate. In Chapter 2, we demonstrated that this distance-based search works for predicting partially observed future states of a network process, but its fundamental bottleneck is the need to enumerate a possibly exponential number of future state candidates. To solve the problem of the efficient exploration of the state space of a network process, we need to utilize the domain knowledge about the structure of this state space, and exploit semi-metricity of the underlying distance measure, aiming at drastically reducing the size of the fraction of the state space that we actually need to traverse.

Combinatorial Network Process Identification: Above, we have advocated the model-based network process analysis approach, having assumed that we have a model for an ob-

served network process. In many cases, we may indeed have a reliable model for the process, as in the case with the network version of the heat equation for the process of heat diffusion in a material network, or Kuramoto model for the process of neuronal synchronization between pairs of neurons in a brain network. However, we may have a process whose nature and, hence, model are unclear, or a process for which there are many alternative models and it is not obvious which model fits the process best. Having observed the evolution of a network process, the key question here is *what is the (best) model for a given process*. While there are methods for identification of non-linear dynamical systems, there are no such methods for combinatorial processes, such as those characterized by combinatorial models of information diffusion through a social network. An extension of a network process identification problem is a problem of disentangling a mixture of processes occurring in the same network simultaneously. Clearly, there is a need for the development of new theories and algorithms that would enable combinatorial network process identification.

5.2 Open Problems in Network Process Modeling

The future of network process modeling research has two primary directions: design of new fundamental modeling tools and accompanying theoretical and algorithmic techniques; and design of models for specific network processes from various domains of human life.

Simpler Invariance Principle for Non-smooth Lyapunov Functions: One potential development of the fundamentals of network process modeling is the extension of the classical (Barabashin-Krasovskiy-)LaSalle Invariance Principle [111, 112]. The standard version of the principle allows to analyze convergence of non-linear dynamical systems, and, similarly to Lyapunov's Second Method [110], requires existence of a smooth Lyapunov function, also known as the energy function, whose behavior along the trajectories of the system allows mak-

ing conclusions about the behavior of the system itself. Finding Lyapunov functions can be rather difficult for systems defined over directed networks (see [212] for a representative result). We addressed the latter issue in Chapter 3 by resorting to the tools from non-smooth analysis. However, the way we exploit the generalized gradient and Lie derivative, and, more specifically, our treatment of the cases when the set-valued Lie derivative is an empty set in our convergence theorems suggests that the statement of LaSalle's Theorem can be generalized and used with non-smooth Lyapunov functions without the need to resort to non-smooth analysis.

New Non-linear Models for Socio-Economic Network Processes: There are great many opportunities for the design of non-linear models for a variety of processes, especially those occurring in social and economic networks. One direction for prospective research is the design of sociologically plausible opinion dynamics models where social ties evolve in time based on the users' positions in the society. For example, a rising politician's opinion can gain weight even outside of his or her network neighborhood solely due to that person's gaining social power. This idea has been explored in the past in the context of DeGroot–Friedkin model for the evolution of social influence networks proposed by Jia et al. [117], and recently revisited by Ye et al. [213]. This model allows the social tie weights to occasionally change based on the users' evolving centrality, but, since centrality is a global measure, the agents are assumed to observe the entire network. One may want to design a non-linear DeGroot-type opinion formation model, where (i) *all social ties* in the network *continuously evolve* based upon the agents' social positions; and (ii) the agents weight their social ties based on *locally available information*—the states and social positions of either immediate neighbors of those about whom one can learn by interacting with neighbors. Such a model would be useful for applications to large-scale online social networks, which are inherently incomplete, and where users make decisions based on a small observed part of the network.

Additional promising research directions for network processes modeling include tractable generalization of Heider-Cartwright-Harary structural balance theory [94, 93] to incomplete social networks, and the design of opinion formation models in complex hierarchical networks, where opinion formation is governed by different laws at different network scales. The latter research may potentially borrow some ideas from multi-layer network theory [214], as well as some ideas from physics, such as the notion of the renormalization group [215].

5.3 Open Problems in Network Process Control

Control Theory for Social, Economic, and Biological Systems: The classic control theory is a well-established field, providing us with the fundamental tools for mathematically characterizing systems, typically, by the means of differential or difference equations, analyzing their behavior, and designing controllers for them. It is, however, not as well understood yet how to control biological and economic systems, while the topic of social control—studied for decades in sociology and political science—hardly has any mathematical or algorithmic foundation. Development of the *social control theory* is critical today, as online social networks have permeated many areas of our lives, exerting implicit influence upon our day-to-day decision making. However, traditional control theory is not immediately transferrable to social systems. For example, controllability—defined as a system’s ability to be stirred to any state by a controller—is meaningless in the context of systems comprised of humans, as in any *socium* there are states unreachable by the controller. The latter calls for quantification of controllability suitable for social systems—that could reflect, for example, the amount of “effort” it would take the controller to stir the system to a desired state—and, more generally, for the creation of a theoretical framework for studying social control and, more importantly, defending against it.

Additional research direction in network process control is the design of collaboration incentives and constraints to optimize performance of collaboration and financial networks—which is closely aligned with the ongoing research in the field of mechanism design—as well as the design of mechanisms to control biological systems, such as brain networks, to fight disease.

5.4 Future of Network Process Research

As the world becomes more interlinked, the network processes will clearly remain in the spotlight. Rapid development of online social network and economic institutions as well as improvement of our understanding of naturally occurring network processes on one side, and evolution of theoretical techniques on the other side allow the studies of network processes to enter a qualitatively new stage. The network process research—conducted along the three facets of analysis, modeling, and control—involves design, theoretical analysis, and optimization of domain-specific plausible models, as well as the development of practical scalable algorithms to make these models work in real-world applications. This research agenda requires bridging domain-specific areas, such as social psychology and behavioral economics, biology, neuroscience, and robotics, as well as the fundamental fields of combinatorial algorithm design, game theory and mechanism design, machine learning, linear algebra, and dynamical systems theory. This thesis, touching upon some of these fields, contributes to and brings closer the formation of the analysis, modeling, and control of dynamic processes in networks as an independent cross-discipline.

Bibliography

- [1] A.-L. Barabási and M. Pósfai, *Network Science*. Cambridge University Press, 2016.
- [2] L. Euler, *Leonhard Euler and the Königsberg bridges*, *Scientific American* **189** (1953), no. 1 66–72.
- [3] L. Katz, *A new status index derived from sociometric analysis*, *Psychometrika* **18** (1953), no. 1 39–43.
- [4] P. Bonacich, *Factoring and weighting approaches to status scores and clique identification*, *Journal of Mathematical Sociology* **2** (1972), no. 1 113–120.
- [5] L. C. Freeman, *A set of measures of centrality based on betweenness*, *Sociometry* (1977) 35–41.
- [6] U. Brandes, *A faster algorithm for betweenness centrality*, *Journal of Mathematical Sociology* **25** (2001), no. 2 163–177.
- [7] N. E. Friedkin, *Theoretical foundations for centrality measures*, *American Journal of Sociology* **96** (1991), no. 6 1478–1504.
- [8] L. Page, S. Brin, R. Motwani, and T. Winograd, *The PageRank citation ranking: Bringing order to the web*, tech. rep., Stanford InfoLab, 1999.
- [9] J. M. Kleinberg, *Authoritative sources in a hyperlinked environment*, *Journal of the ACM* **46** (1999), no. 5 604–632.
- [10] M. E. Newman, *A measure of betweenness centrality based on random walks*, *Social Networks* **27** (2005), no. 1 39–54.
- [11] P. W. Holland and S. Leinhardt, *Transitivity in structural models of small groups*, *Comparative Group Studies* **2** (1971), no. 2 107–124.
- [12] D. J. Watts and S. H. Strogatz, *Collective dynamics of “small-world” networks*, *Nature* **393** (1998), no. 6684 440.
- [13] J. Saramäki, M. Kivelä, J.-P. Onnela, K. Kaski, and J. Kertesz, *Generalizations of the clustering coefficient to weighted complex networks*, *Physical Review E* **75** (2007), no. 2 027105.

- [14] M. R. Garey and D. S. Johnson, *Computers and Intractability*, vol. 29. W.H. Freeman, New York, 2002.
- [15] F. R. Chung and F. C. Graham, *Spectral Graph Theory*. American Mathematical Society, 1997.
- [16] M. Fiedler, *Algebraic connectivity of graphs*, *Czechoslovak Mathematical Journal* **23** (1973), no. 2 298–305.
- [17] J. Shi and J. Malik, *Normalized cuts and image segmentation*, *IEEE Trans. on Pattern Analysis and Machine Intelligence* **22** (2000), no. 8 888–905.
- [18] A. Y. Ng, M. I. Jordan, and Y. Weiss, *On spectral clustering: Analysis and an algorithm*, in *Advances in Neural Information Processing Systems*, pp. 849–856, 2002.
- [19] U. Von Luxburg, *A tutorial on spectral clustering*, *Statistics and Computing* **17** (2007), no. 4 395–416.
- [20] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney, *Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters*, *Internet Mathematics* **6** (2009), no. 1 29–123.
- [21] S. Fortunato, *Community detection in graphs*, *Physics Reports* **486** (2010), no. 3-5 75–174.
- [22] V. Hodge and J. Austin, *A survey of outlier detection methodologies*, *Artificial Intelligence Review* **22** (2004), no. 2 85–126.
- [23] X. Yan, P. S. Yu, and J. Han, *Substructure similarity search in graph databases*, in *Proc. of Special Interest Group on Management of Data (SIGMOD)*, pp. 766–777, ACM, 2005.
- [24] X. Yan and J. Han, *gSpan: Graph-based substructure pattern mining*, in *Proc. of International Conf. on Data Mining (ICDM)*, pp. 721–724, IEEE, 2002.
- [25] C. C. Aggarwal and J. Han, *Frequent Pattern Mining*. Springer, 2014.
- [26] P. Turán, *On an external problem in graph theory*, *Mat. Fiz. Lapok* **48** (1941) 436–452.
- [27] E. W. Dijkstra, *A note on two problems in connexion with graphs*, *Numerische Mathematik* **1** (1959), no. 1 269–271.
- [28] D. B. Johnson, *Efficient algorithms for shortest paths in sparse networks*, *Journal of the ACM* **24** (1977), no. 1 1–13.
- [29] M. L. Fredman and R. E. Tarjan, *Fibonacci heaps and their uses in improved network optimization algorithms*, *Journal of the ACM (JACM)* **34** (1987), no. 3 596–615.

- [30] R. K. Ahuja, K. Mehlhorn, J. Orlin, and R. E. Tarjan, *Faster algorithms for the shortest path problem*, *Journal of the ACM* **37** (1990), no. 2 213–223.
- [31] L. R. Ford and D. R. Fulkerson, *Maximal flow through a network*, *Canadian Journal of Mathematics* **8** (1956), no. 3 399–404.
- [32] J. Edmonds and R. M. Karp, *Theoretical improvements in algorithmic efficiency for network flow problems*, *Journal of the ACM (JACM)* **19** (1972), no. 2 248–264.
- [33] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows*. Elsevier, 2014.
- [34] J. Bourgain, *On Lipschitz embedding of finite metric spaces in Hilbert space*, *Israel Journal of Mathematics* **52** (1985), no. 1-2 46–52.
- [35] N. Linial, E. London, and Y. Rabinovich, *The geometry of graphs and some of its algorithmic applications*, *Combinatorica* **15** (1995), no. 2 215–245.
- [36] D. Conte, P. Foggia, C. Sansone, and M. Vento, *Thirty years of graph matching in pattern recognition*, *International Journal of Pattern Recognition and Artificial Intelligence* **18** (2004), no. 03 265–298.
- [37] K. Riesen and H. Bunke, *Graph classification by means of lipschitz embedding*, *IEEE Trans. on Systems, Man, and Cybernetics, Part B: Cybernetics* **39** (2009), no. 6 1472–1483.
- [38] K. Riesen and H. Bunke, *Graph classification and clustering based on vector space embedding*, vol. 77. World Scientific, 2010.
- [39] K. M. Borgwardt, H.-P. Kriegel, and P. Wackersreuther, *Pattern mining in frequent dynamic subgraphs*, in *Proc. of International Conf. on Data Mining (ICDM)*, pp. 818–822, IEEE, 2006.
- [40] A. Silva, A. Singh, and A. Swami, *Spectral algorithms for temporal graph cuts*, in *Proc. of International Conf. on World Wide Web (WWW)*, pp. 519–528, 2018.
- [41] K. Lerman, R. Ghosh, and J. H. Kang, *Centrality metric for dynamic networks*, in *Proc. of ACM Workshop on Mining and Learning with Graphs*, pp. 70–77, 2010.
- [42] H. Kim and R. Anderson, *Temporal node centrality in complex networks*, *Physical Review E* **85** (2012), no. 2 026107.
- [43] D. Taylor, S. A. Myers, A. Clauset, M. A. Porter, and P. J. Mucha, *Eigenvector-based centrality measures for temporal networks*, *Multiscale Modeling & Simulation* **15** (2017), no. 1 537–574.
- [44] P. Erdős and A. Rényi, *On random graphs I*, *Publicationes Mathematicae (Debrecen)* **6** (1959) 290–297.

- [45] R. Albert and A.-L. Barabási, *Statistical mechanics of complex networks*, *Reviews of Modern Physics* **74** (2002), no. 1 47.
- [46] M. O. Jackson, *Social and Economic Networks*. Princeton University Press, 2010.
- [47] A. Barrat, M. Barthelemy, and A. Vespignani, *Dynamical Processes on Complex Networks*. Cambridge University Press, 2008.
- [48] M. Gardner, *Mathematical games – The fantastic combinations of John Conway’s new solitaire game “life”*, *Scientific American* **223** (1970) 120–123.
- [49] R. P. Abelson, *Mathematical models of the distribution of attitudes under controversy*, *Contributions to Mathematical Psychology* **14** (1964) 1–160.
- [50] C. Tan, L. Lee, J. Tang, L. Jiang, M. Zhou, and P. Li, *User-level sentiment analysis incorporating social networks*, in *Proc. of International Conf. on Knowledge Discovery and Data Mining (KDD)*, pp. 1397–1405, ACM, 2011.
- [51] X. Wang, F. Wei, X. Liu, M. Zhou, and M. Zhang, *Topic sentiment analysis in Twitter: A graph-based hashtag sentiment classification approach*, in *Proc. of International Conf. on Information and Knowledge Management (CIKM)*, pp. 1031–1040, ACM, 2011.
- [52] T. Carnes, C. Nagarajan, S. M. Wild, and A. Van Zuylen, *Maximizing influence in a competitive social network: A follower’s perspective*, in *Proc. of Electronic Commerce*, pp. 351–360, ACM, 2007.
- [53] A. Goyal, F. Bonchi, and L. V. Lakshmanan, *Learning influence probabilities in social networks*, in *Proc. of International Conf. on Web Search and Data Mining (WSDM)*, pp. 241–250, ACM, 2010.
- [54] A. Borodin, Y. Filmus, and J. Oren, *Threshold models for competitive influence in social networks*, in *Proc. of International Workshop on Internet and Network Economics*, pp. 539–550. Springer, 2010.
- [55] Y. Rubner, C. Tomasi, and L. J. Guibas, *The Earth Mover’s Distance as a metric for image retrieval*, *International Journal of Computer Vision* **40** (2000), no. 2 99–121.
- [56] K. L. Clarkson, *Nearest-neighbor searching and metric space dimensions*, in *Nearest-Neighbor Methods in Learning and Vision: Theory and Practice* (G. Shakhnarovich, T. Darrell, and P. Indyk, eds.). The MIT Press, 2006.
- [57] J. Hafner, H. S. Sawhney, W. Equitz, M. Flickner, and W. Niblack, *Efficient color histogram indexing for quadratic form distance functions*, *Pattern Analysis and Machine Intelligence* **17** (1995), no. 7 729–736.

- [58] H. Bunke and K. Shearer, *A graph distance metric based on the maximal common subgraph*, *Pattern Recognition Letters* **19** (1998), no. 3 255–259.
- [59] X. Gao, B. Xiao, D. Tao, and X. Li, *A survey of Graph Edit Distance*, *Pattern Analysis and Applications* **13** (2010), no. 1 113–129.
- [60] D. Koutra, J. T. Vogelstein, and C. Faloutsos, *DeltaCon: A principled massive-graph similarity function*, in *Proc. of International Conf. on Data Mining (SDM)*, pp. 162–170, SIAM, 2013.
- [61] V. D. Blondel, A. Gajardo, M. Heymans, P. Senellart, and P. Van Dooren, *A measure of similarity between graph vertices: Applications to synonym extraction and web searching*, *SIAM Review* **46** (2004), no. 4 647–666.
- [62] E. A. Leicht, P. Holme, and M. E. J. Newman, *Vertex similarity in networks*, *Physical Review E* **73** (2006), no. 2 026120.
- [63] S. Melnik, H. Garcia-Molina, and E. Rahm, *Similarity flooding: A versatile graph matching algorithm and its application to schema matching*, in *Proc. of International Conf. on Data Engineering (ICDE)*, pp. 117–128, IEEE, 2002.
- [64] G. Jeh and J. Widom, *SimRank: A measure of structural-context similarity*, in *Proc. of International Conf. on Knowledge Discovery and Data Mining (KDD)*, pp. 538–543, ACM, 2002.
- [65] S. Segarra, W. Huang, and A. Ribeiro, *Diffusion and superposition distances for signals supported on networks*, *IEEE Signal and Information Processing over Networks* **1** (2015), no. 1 20–32.
- [66] D. K. Hammond, Y. Gur, and C. R. Johnson, *Graph diffusion distance: A difference measure for weighted graphs based on the graph Laplacian exponential kernel*, in *Proc. of Global Conference on Signal and Information Processing*, pp. 419–422, 2013.
- [67] H. Ling and K. Okada, *Diffusion distance for histogram comparison*, in *Proc. of Computer Vision and Pattern Recognition*, vol. 1, pp. 246–253, 2006.
- [68] M. Berlingerio, D. Koutra, T. Eliassi-Rad, and C. Faloutsos, *Network similarity via multiple social theories*, in *Proc. of International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pp. 1439–1440, IEEE, 2013.
- [69] R. C. Wilson, E. R. Hancock, and B. Luo, *Pattern vectors from algebraic graph theory*, *IEEE Pattern Analysis and Machine Intelligence* **27** (2005), no. 7 1112–1124.
- [70] P. Zhu and R. C. Wilson, *A study of graph spectra for comparing graphs*, in *Proc. of BMVA British Machine Vision Conference*, 2005.

- [71] M. Gomez-Rodriguez, J. Leskovec, and A. Krause, *Inferring networks of diffusion and influence*, in *Proc. of International Conf. on Knowledge Discovery and Data Mining (KDD)*, pp. 1019–1028, ACM, 2010.
- [72] O. Pele and M. Werman, *A linear time histogram metric for improved SIFT matching*, in *Proc. of European Conference on Computer Vision*, pp. 495–508. Springer, 2008.
- [73] V. Ljosa, A. Bhattacharya, and A. K. Singh, *Indexing spatially sensitive distance measures using multi-resolution lower bounds*, *Proc. of International Conference on Extending Database Technology* (2006).
- [74] N. Karmarkar, *A new polynomial-time algorithm for linear programming*, in *Proc. of Symposium on Theory of Computing (STOC)*, pp. 302–311, ACM, 1984.
- [75] F. S. Hillier and G. J. Lieberman, *Introduction to Mathematical Programming*, 1995.
- [76] Y. Tang, U. Leong Hou, Y. Cai, N. Mamoulis, and R. Cheng, *Earth Mover’s Distance based similarity search at scale*, *Trans. of International Journal on Very Large Data Bases (VLDB)* **7** (2013), no. 4 313–324.
- [77] L. Li, M. Ma, P. Lei, X. Wang, and X. Chen, *A linear approximate algorithm for Earth Mover’s Distance with thresholded ground distance*, *Mathematical Problems in Engineering* (2014).
- [78] A. McGregor and D. Stubbs, *Sketching Earth Mover’s Distance on graph metrics*, in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques* (M. Goemans, K. Jansen, J. Rolim, and L. Trevisan, eds.), pp. 274–286. Springer, 2013.
- [79] A. Goldberg and R. Tarjan, *Solving minimum-cost flow problems by successive approximation*, in *Proc. of Symposium on Theory of Computing (STOC)*, pp. 7–18, ACM, 1987.
- [80] R. K. Ahuja, J. B. Orlin, C. Stein, and R. E. Tarjan, *Improved algorithms for bipartite network flow*, *SIAM Journal on Computing* **23** (1994), no. 5 906–933.
- [81] K. Macropol, P. Bogdanov, A. K. Singh, L. Petzold, and X. Yan, *I act, therefore I judge: Network sentiment dynamics based on user activity change*, *Proc. of International Conference on Advances in Social Networks Analysis and Mining (ASONAM)* (2013) 396–402.
- [82] K. Macropol, P. Bogdanov, A. K. Singh, L. Petzold, and X. Yan, “I act, therefore I judge: Network sentiment dynamics based on user activity change (Supplemental Material).” <http://cs.ucsb.edu/~dbl/papers/sentimentappendix.pdf>, 2013. [accessed on 2018-06-15].

- [83] R. Merris, *Laplacian matrices of graphs: A survey*, *Linear Algebra and Its Applications* **197** (1994) 143–176.
- [84] A. Najar, L. Denoyer, and P. Gallinari, *Predicting information diffusion on social networks with partial knowledge*, in *Proc. International Conference on World Wide Web Companion (WWW Companion)* (L. Car, ed.), pp. 1197–1204, 2012.
- [85] K. Saito, M. Kimura, K. Ohara, and H. Motoda, *Detecting changes in opinion value distribution for voter model*, in *Proc. of Social Computing, Behavioral-Cultural Modeling and Prediction*, pp. 89–96, Springer, 2011.
- [86] M. D. Conover, B. Gonçalves, J. Ratkiewicz, A. Flammini, and F. Menczer, *Predicting the political alignment of Twitter users*, in *Proc. of International Conference on Social Computing and Networking (SocialCom)*, pp. 192–199, IEEE, 2011.
- [87] R. Cohen and D. Ruths, *Classifying Political Orientation on Twitter: It's Not Easy!*, in *Proc. of International Conference on Web and Social Media (ICWSM)*, AAAI, 2013.
- [88] A. Goldberg, *An efficient implementation of a scaling minimum-cost flow algorithm*, *Journal of Algorithms* **22** (1997), no. 1 1–29.
- [89] CPLEX, *V12. 1: User's Manual for CPLEX*, International Business Machines Corp. **46** (2009), no. 53 157.
- [90] L. Festinger, *A theory of social comparison processes*, *Human Relations* **7** (1954), no. 2 117–140.
- [91] M. H. DeGroot, *Reaching a consensus*, *Journal of the American Statistical Association* **69** (1974), no. 345 118–121.
- [92] L. Festinger, *A Theory of Cognitive Dissonance*, vol. 2. Stanford University Press, 1962.
- [93] D. Cartwright and F. Harary, *Structural balance: A generalization of Heider's theory*, *Psychological Review* **63** (1956), no. 5 277.
- [94] F. Heider, *Attitudes and cognitive organization*, *The Journal of Psychology* **21** (1946), no. 1 107–112.
- [95] E. M. Pomerantz, S. Chaiken, and R. S. Tordesillas, *Attitude strength and resistance processes*, *Journal of Personality and Social Psychology* **69** (1995), no. 3 408–419.
- [96] R. P. Abelson, *Conviction*, *American Psychologist* **43** (1988), no. 4 267.
- [97] J. A. Krosnick, D. S. Boninger, Y. C. Chuang, M. K. Berent, and C. G. Carnot, *Attitude strength: One construct or many related constructs?*, *Journal of Personality and Social Psychology* **65** (1993), no. 6 1132–1151.

- [98] N. E. Friedkin and E. C. Johnsen, *Social influence networks and opinion change*, *Advances in Group Processes* **16** (1999), no. 1 1–29.
- [99] D. Krackhardt, *A plunge into networks*, *Science* **326** (2009), no. 5949 47–48.
- [100] J. L. Lastovicka and D. M. Gardner, *Components of involvement*, in *Attitude Research Plays for High Stakes* (J. C. Maloney and B. Silverman, eds.), pp. 53–73. American Marketing Association, 1979.
- [101] N. E. Friedkin, *The problem of social control and coordination of complex systems in sociology: A look at the community cleavage problem*, *IEEE Control Systems Magazine* **35** (June, 2015) 40–51.
- [102] J. A. Krosnick and R. E. Petty, *Attitude strength: An overview*, in *Attitude Strength: Antecedents and Consequences*, ch. 1, pp. 1–24. Lawrence Erlbaum Associates, Mahwah, NJ, 1995.
- [103] J. N. Bassili, *Meta-judgmental versus operative indexes of psychological attributes: The case of measures of attitude strength*, *Journal of Personality and Social Psychology* **71** (1996), no. 4 637.
- [104] C. E. Osgood and P. H. Tannenbaum, *The principle of congruity in the prediction of attitude change*, *Psychological Review* **62** (January, 1955) 42–55.
- [105] A. G. Miller, J. W. McHoskey, C. M. Bane, and T. G. Dowd, *The attitude polarization phenomenon: Role of response measure, attitude extremity, and behavioral consequences of reported attitude change*, *Journal of Personality and Social Psychology* **64** (1993), no. 4 561.
- [106] A. H. Eagly and S. Chaiken, *The Psychology of Attitudes*. Harcourt Brace Jovanovich College Publishers, 1993.
- [107] J. N. Bassili and R. D. Brown, *Implicit and explicit attitudes: Research, challenges, and theory*, in *Handbook of Attitudes and Attitude Change* (D. Albarracín, B. T. Johnson, and M. P. Zanna, eds.), ch. 13, pp. 543–574. Lawrence Erlbaum Associates, 2005.
- [108] N. Schwarz and G. Bohner, *The construction of attitudes*, in *Blackwell Handbook of Social Psychology: Intraindividual Processes* (A. Tesser and N. Schwarz, eds.), ch. 20, pp. 436–457. Blackwell Publishing, 2001.
- [109] J. Zaller, *The Nature and Origins of Mass Opinion*. Cambridge University Press, 1992.
- [110] A. M. Lyapunov, *The general problem of the stability of motion*, *International Journal of Control* **55** (1892), no. 3 531–534.
- [111] E. Barbashin and N. N. Krasovskiy, *On stability of motion in the large*, in *Reports of the USSR Academy of Science*, vol. 86(3), pp. 453–456, 1952.

- [112] J. LaSalle, *Some extensions of Lyapunov's second method*, *IRE Trans. on Circuit Theory* **7** (1960), no. 4 520–527.
- [113] L. Moreau, *Stability of continuous-time distributed consensus algorithms*, in *IEEE Conf. on Decision and Control*, vol. 4, (Nassau, Bahamas), pp. 3998–4003, 2004.
- [114] J. R. French Jr., *A formal theory of social power*, *Psychological Review* **63** (1956), no. 3 181–194.
- [115] F. Harary, *A criterion for unanimity in French's theory of social power*, in *Studies in Social Power* (D. Cartwright, ed.), pp. 168–182. University of Michigan, 1959.
- [116] M. Pirani and S. Sundaram, *Spectral properties of the grounded Laplacian matrix with applications to consensus in the presence of stubborn agents*, in *American Control Conference*, pp. 2160–2165, IEEE, 2014.
- [117] P. Jia, A. MirTabatabaei, N. E. Friedkin, and F. Bullo, *Opinion dynamics and the evolution of social power in influence networks*, *SIAM Review* **57** (2015), no. 3 367–397.
- [118] Z. Xu, J. Liu, and T. Başar, *On a modified DeGroot-Friedkin model of opinion dynamics*, in *American Control Conference*, pp. 1047–1052, IEEE, July, 2015.
- [119] P. Groeber, J. Lorenz, and F. Schweitzer, *Dissonance minimization as a microfoundation of social influence in models of opinion formation*, *The Journal of Mathematical Sociology* **38** (2014), no. 3 147–174.
- [120] J. Ghaderi and R. Srikant, *Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate*, *Automatica* **50** (2014), no. 12 3209–3215.
- [121] C. Castellano, S. Fortunato, and V. Loreto, *Statistical physics of social dynamics*, *Reviews of Modern Physics* **81** (May, 2009) 591.
- [122] J. Lorenz, *Continuous opinion dynamics under bounded confidence: A survey*, *International Journal of Modern Physics C* **18** (2007), no. 12 1819–1838.
- [123] R. Hegselmann and U. Krause, *Opinion dynamics and bounded confidence models, analysis, and simulation*, *Journal of Artificial Societies and Social Simulation* **5** (2002), no. 3.
- [124] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, *Mixing beliefs among interacting agents*, *Advances in Complex Systems* **3** (2000), no. 1/4 87–98.
- [125] V. D. Blondel, J. M. Hendrickx, and J. N. Tsitsiklis, *On Krause's multi-agent consensus model with state-dependent connectivity*, *IEEE Trans. on Automatic Control* **54** (2009), no. 11 2586–2597.

- [126] A. MirTabatabaei and F. Bullo, *Opinion dynamics in heterogeneous networks: Convergence conjectures and theorems*, *SIAM Journal on Control and Optimization* **50** (2012), no. 5 2763–2785.
- [127] J. Lorenz, *A stabilization theorem for dynamics of continuous opinions*, *Physica A: Statistical Mechanics and its Applications* **355** (2005), no. 1 217–223.
- [128] G. Deffuant, F. Amblard, and G. Weisbuch, *Modelling group opinion shift to extreme: The smooth bounded confidence model*, *arXiv:cond-mat/0410199 [cond-mat.other]* (2004).
- [129] E. Kurmyshev, H. A. Juárez, and R. A. González-Silva, *Dynamics of bounded confidence opinion in heterogeneous social networks: Concord against partial antagonism*, *Physica A: Statistical Mechanics and its Applications* **390** (2011), no. 16 2945–2955.
- [130] M. Jalili, *Effects of leaders and social power on opinion formation in complex networks*, *Simulation* **89** (2013), no. 5 578–588.
- [131] P. Sobkowicz, *Extremism without extremists: Deffuant model with emotions*, *Frontiers in Physics* **3** (2015) 17.
- [132] S. Chen, D. H. Glass, and M. McCartney, *Characteristics of successful opinion leaders in a bounded confidence model*, *Physica A: Statistical Mechanics and its Applications* **449** (2016) 426–436.
- [133] K. Tucci, J. González-Avella, and M. Cosenza, *Rise of an alternative majority against opinion leaders*, *Physica A: Statistical Mechanics and its Applications* **446** (2016) 75–81.
- [134] Y. Cao, W. Yu, W. Ren, and G. Chen, *An overview of recent progress in the study of distributed multi-agent coordination*, *IEEE Trans. on Industrial Informatics* **9** (2013), no. 1 427–438.
- [135] R. Olfati-Saber, A. Fax, and R. M. Murray, *Consensus and cooperation in networked multi-agent systems*, *Proc. of the IEEE* **95** (2007), no. 1 215–233.
- [136] S. Martin and J. Hendrickx, *Continuous-time consensus under non-instantaneous reciprocity*, *IEEE Trans. on Automatic Control* **PP** (2015), no. 99.
- [137] L. Moreau, *Stability of multiagent systems with time-dependent communication links*, *IEEE Trans. on Automatic Control* **50** (2005), no. 2 169–182.
- [138] R. Olfati-Saber and R. M. Murray, *Consensus problems in networks of agents with switching topology and time-delays*, *IEEE Trans. on Automatic Control* **49** (2004), no. 9 1520–1533.

- [139] C. Altafini, *Consensus problems on networks with antagonistic interactions*, *IEEE Trans. on Automatic Control* **58** (2013), no. 4 935–946.
- [140] J. M. Hendrickx, *A lifting approach to models of opinion dynamics with antagonisms*, in *IEEE Conf. on Decision and Control*, pp. 2118–2123, 2014.
- [141] F. Fagnani and S. Zampieri, *Randomized consensus algorithms over large scale networks*, *IEEE Journal on Selected Areas in Communications* **26** (2008), no. 4 634–649.
- [142] A. Tahbaz-Salehi and A. Jadbabaie, *A necessary and sufficient condition for consensus over random networks*, *IEEE Trans. on Automatic Control* **53** (2008), no. 3 791–795.
- [143] A. Waagen, G. Verma, K. Chan, A. Swami, and R. D’Souza, *Effect of zealotry in high-dimensional opinion dynamics models*, *Physical Review E* **91** (2015), no. 2 022811.
- [144] J. Goldenberg, B. Libai, and E. Muller, *Talk of the network: A complex systems look at the underlying process of word-of-mouth*, *Marketing Letters* **12** (2001), no. 3 211–223.
- [145] D. Kempe, J. Kleinberg, and É. Tardos, *Maximizing the spread of influence through a social network*, in *Proc. of International Conf. on Knowledge Discovery and Data Mining (KDD)*, pp. 137–146, ACM, 2003.
- [146] K. Saito, M. Kimura, K. Ohara, and H. Motoda, *Selecting information diffusion models over social networks for behavioral analysis*, in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 180–195. Springer, 2010.
- [147] S. Goyal and M. Kearns, *Competitive contagion in networks*, in *Proc. of Symposium on Theory of Computing (STOC)*, pp. 759–774, ACM, 2012.
- [148] P. Clifford and A. Sudbury, *A model for spatial conflict*, *Biometrika* **60** (1973), no. 3 581–588.
- [149] E. Even-Dar and A. Shapira, *A note on maximizing the spread of influence in social networks*, in *Proc. of International Workshop on Internet and Network Economics*, pp. 281–286. Springer, 2007.
- [150] M. Kimura, K. Saito, K. Ohara, and H. Motoda, *Learning to predict opinion share in social networks.*, in *AAAI*, 2010.
- [151] E. Yildiz, D. Acemoglu, A. Ozdaglar, A. Saberi, and A. Scaglione, *Discrete opinion dynamics with stubborn agents*, *Operations Research, OPRE-2011-01-026* (2011).
- [152] M. Granovetter, *Threshold models of collective behavior*, *American Journal of Sociology* (1978) 1420–1443.

- [153] D. J. Watts and P. S. Dodds, *Influentials, networks, and public opinion formation*, *Journal of Consumer Research* **34** (2007), no. 4 441–458.
- [154] D. Acemoglu, G. Como, F. Fagnani, and A. Ozdaglar, *Opinion fluctuations and disagreement in social networks*, *Mathematics of Operations Research* **38** (2013), no. 1 1–27.
- [155] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione, *Binary opinion dynamics with stubborn agents*, *ACM Trans. on Economics and Computation* **1** (2013), no. 4 19.
- [156] D. Acemoglu and A. Ozdaglar, *Opinion dynamics and learning in social networks*, *Dynamic Games and Applications* **1** (2011), no. 1 3–49.
- [157] A. G. Whittington, A. M. Hofmeister, and P. I. Nabelek, *Temperature-dependent thermal diffusivity of the Earth’s crust and implications for magmatism*, *Nature* **458** (2009), no. 7236 319–321.
- [158] M. Taylor, *Towards a mathematical theory of influence and attitude change*, *Human Relations* **21** (1968), no. 2 121–139.
- [159] F. H. Clarke, *Optimization and Nonsmooth Analysis*, vol. 5 of *Classics in Applied Mathematics*. SIAM, 1990.
- [160] J. Cortés, *Discontinuous dynamical systems – A tutorial on solutions, nonsmooth analysis, and stability*, *IEEE Control Systems Magazine* **28** (June, 2008) 36–73.
- [161] A. Bacciotti and F. Ceragioli, *Stability and stabilization of discontinuous systems and nonsmooth Lyapunov functions*, *ESAIM: Control, Optimisation & Calculus of Variations* **4** (1999) 361–376.
- [162] H. K. Khalil, *Nonlinear Systems*. Prentice Hall, 3 ed., 2002.
- [163] Z. Lin, B. Francis, and M. Maggiore, *State agreement for continuous-time coupled nonlinear systems*, *SIAM Journal on Control and Optimization* **46** (2007), no. 1 288–307.
- [164] W. Zachary, *An information flow model for conflict and fission in small groups*, *Journal of Anthropological Research* **33** (1977), no. 4 452–473.
- [165] J. M. Hendrickx, *Graphs and networks for the analysis of autonomous agent systems*. PhD thesis, Université Catholique de Louvain, Belgium, February, 2008.
- [166] V. Amelkin, F. Bullo, and A. K. Singh, *Polar opinion dynamics in social networks*, *IEEE Trans. on Automatic Control* **62** (Nov., 2017).

- [167] F. Bullo, *Lectures on Network Systems*. CreateSpace, 1 ed., 2018. With contributions by J. Cortes, F. Dorfler, and S. Martinez.
- [168] J. G. Kemeny and J. L. Snell, *Finite Markov Chains*. Springer, 1976.
- [169] V. Chandola, A. Banerjee, and V. Kumar, *Anomaly detection for discrete sequences: A survey*, *IEEE Trans. on Knowledge and Data Engineering (TKDE)* **24** (2012), no. 5 823–839.
- [170] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance, *Cost-effective outbreak detection in networks*, in *Proc. of International Conf. on Knowledge Discovery and Data Mining (KDD)*, pp. 420–429, ACM, 2007.
- [171] V. Amelkin, P. Bogdanov, and A. K. Singh, *A distance measure for the analysis of polar opinion dynamics in social networks*, in *Proc. of International Conf. on Data Engineering (ICDE)*, pp. 159–162, IEEE, 2017.
- [172] R. M. Karp, *Reducibility among combinatorial problems*, in *Complexity of Computer Computations*, pp. 85–103. Springer, 1972.
- [173] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, *An analysis of approximations for maximizing submodular set functions—I*, *Mathematical Programming* **14** (1978), no. 1 265–294.
- [174] A. Ghosh and S. Boyd, *Growing well-connected graphs*, in *IEEE Conf. on Decision and Control*, pp. 6605–6611, 2006.
- [175] Z. Yu, C. Wang, J. Bu, X. Wang, Y. Wu, and C. Chen, *Friend recommendation with content spread enhancement in social networks*, *Information Sciences* **309** (2015) 102–118.
- [176] P. Van Mieghem, D. Stevanović, F. Kuipers, C. Li, R. Van De Bovenkamp, D. Liu, and H. Wang, *Decreasing the spectral radius of a graph by link removals*, *Physical Review E* **84** (2011), no. 1 016101.
- [177] S. Saha, A. Adiga, B. A. Prakash, and A. K. S. Vullikanti, *Approximation algorithms for reducing the spectral radius to control epidemic spread*, in *Proc. of International Conf. on Data Mining (SDM)*, pp. 568–576, SIAM, 2015.
- [178] Y. Zhang, A. Adiga, A. Vullikanti, and B. A. Prakash, *Controlling propagation at group scale on networks*, in *Proc. of International Conf. on Data Mining (ICDM)*, pp. 619–628, IEEE, 2015.
- [179] H. Tong, A. Prakash, T. Eliassi-Rad, M. Faloutsos, and C. Faloutsos, *Gelling and melting large graphs by edge manipulation*, in *Proc. of Intl. Conf. on Information and Knowledge Management (CIKM)*, pp. 245–254, ACM, 2012.

- [180] C. Chen, H. Tong, B. A. Prakash, T. Eliassi-Rad, M. Faloutsos, and C. Faloutsos, *Eigen-optimization on large graphs by edge manipulation*, *ACM Trans. on Knowledge Discovery from Data (TKDD)* **10** (2016), no. 4 49.
- [181] L. T. Le, T. Eliassi-Rad, and H. Tong, *MET: A fast algorithm for minimizing propagation in large graphs with small eigen-gaps*, in *Proc. of International Conf. on Data Mining (SDM)*, pp. 694–702, SIAM, 2015.
- [182] H. Chan, L. Akoglu, and H. Tong, *Make it or break it: Manipulating robustness in large networks*, in *Proc. of Intl. Conf. on Data Mining (SDM)*, pp. 325–333, SIAM, 2014.
- [183] A. Ghosh, S. Boyd, and A. Saberi, *Minimizing effective resistance of a graph*, *SIAM Review* **50** (2008), no. 1 37–66.
- [184] F. Arrigo and M. Benzi, *Updating and downdating techniques for optimizing network communicability*, *SIAM Journal on Scientific Computing* **38** (2016), no. 1 25–49.
- [185] K. Garimella, G. De Francisci Morales, A. Gionis, and M. Mathioudakis, *Reducing controversy by connecting opposing views*, in *Proc. of International Conf. on Web Search and Data Mining (WSDM)*, pp. 81–90, ACM, 2017.
- [186] V. Chaoji, S. Ranu, R. Rastogi, and R. Bhatt, *Recommendations to boost content spread in social networks*, in *Proc. of International Conf. on World Wide Web (WWW)*, pp. 529–538, ACM, 2012.
- [187] C. J. Kuhlman, G. Tuli, S. Swarup, M. V. Marathe, and S. Ravi, *Blocking simple and complex contagion by edge removal*, in *Proc. of International Conf. on Data Mining (ICDM)*, pp. 399–408, IEEE, 2013.
- [188] E. B. Khalil, B. Dilkina, and L. Song, *Scalable diffusion-aware optimization of network topology*, in *Proc. of International Conf. on Knowledge Discovery and Data Mining (KDD)*, pp. 1226–1235, ACM, 2014.
- [189] C. A. Phillips, *The network inhibition problem*, in *Proc. of Symposium on Theory of Computing (STOC)*, pp. 776–785, ACM, 1993.
- [190] E. Israeli and R. K. Wood, *Shortest-path network interdiction*, *Networks* **40** (2002), no. 2 97–111.
- [191] M. Papagelis, F. Bonchi, and A. Gionis, *Suggesting ghost edges for a smaller world*, in *Proc. of International Conf. on Information and Knowledge Management (CIKM)*, pp. 2305–2308, ACM, 2011.
- [192] V. Ishakian, D. Erdős, E. Terzi, and A. Bestavros, *A framework for the evaluation and management of network centrality*, in *Proc. of International Conf. on Data Mining (SDM)*, pp. 427–438, SIAM, 2012.

- [193] N. Parotsidis, E. Pitoura, and P. Tsaparas, *Centrality-aware link recommendations*, in *Proc. of International Conf. on Web Search and Data Mining (WSDM)*, pp. 503–512, ACM, 2016.
- [194] S. Perumal, P. Basu, and Z. Guan, *Minimizing eccentricity in composite networks via constrained edge additions*, in *Proc. of Military Communications Conference (MILCOM)*, pp. 1894–1899, IEEE, 2013.
- [195] S. Medya, A. Silva, A. Singh, P. Basu, and A. Swami, *Group centrality maximization via network design*, in *Proc. of International Conf. on Data Mining (SDM)*, pp. 126–134, 2018.
- [196] K. Avrachenkov and N. Litvak, *The effect of new links on Google PageRank*, *Stochastic Models* **22** (2006), no. 2 319–331.
- [197] C. de Kerchove, L. Ninove, and P. Van Dooren, *Maximizing PageRank via outlinks*, *Linear Algebra and Its Applications* **429** (2008), no. 5-6 1254–1276.
- [198] B. C. Csáji, R. M. Jungers, and V. D. Blondel, *PageRank optimization by edge selection*, *Discrete Applied Mathematics* **169** (2014) 73–87.
- [199] G. E. Cho and C. D. Meyer, *Markov chain sensitivity measured by mean first passage times*, *Linear Algebra and Its Applications* **316** (2000), no. 1-3 21–28.
- [200] S. Chien, C. Dwork, R. Kumar, and D. Sivakumar, *Towards exploiting link evolution*, in *Workshop on Algorithms and Models for the Web Graph*, 2001.
- [201] A. N. Langville and C. D. Meyer, *Updating Markov chains with an eye on Google’s PageRank*, *SIAM Journal on Matrix Analysis and Applications* **27** (2006), no. 4 968–987.
- [202] J. J. Hunter, *Stationary distributions and mean first passage times of perturbed Markov chains*, *Linear Algebra and Its Applications* **410** (2005), no. 1-3 217–243.
- [203] G. Como and F. Fagnani, *Robustness of large-scale stochastic matrices to localized perturbations*, *IEEE Trans. on Network Science and Engineering* **2** (2015), no. 2 53–64.
- [204] B. Bahmani, R. Kumar, M. Mahdian, and E. Upfal, *PageRank on an evolving graph*, in *Proc. of Conf. on Knowl. Discovery and Data Mining (KDD)*, pp. 24–32, ACM, 2012.
- [205] L. Li, H. Tong, Y. Xiao, and W. Fan, *Cheetah: Fast graph kernel tracking on dynamic graphs*, in *Proc. of International Conf. on Data Mining (SDM)*, pp. 280–288, SIAM, 2015.
- [206] C. Chen and H. Tong, *On the eigen-functions of dynamic graphs: Fast tracking and attribution algorithms*, *Statistical Analysis and Data Mining: The ASA Data Science Journal* **10** (2017), no. 2 121–135.

- [207] J. J. Hunter, *The computation of the mean first passage times for Markov chains*, *arXiv:1701.07781v2 [math.NA]* (2017).
- [208] S. White and P. Smyth, *Algorithms for estimating relative importance in networks*, in *Proc. of Conf. on Knowl. Discovery and Data Mining (KDD)*, pp. 266–275, ACM, 2003.
- [209] K. Avrachenkov, N. Litvak, D. Nemirovsky, and N. Osipova, *Monte Carlo methods in PageRank computation: When one iteration is sufficient*, *SIAM Journal on Numerical Analysis* **45** (2007), no. 2 890–904.
- [210] J. J. McAuley and J. Leskovec, *Learning to discover social circles in ego networks*, in *Advances in Neural Information Processing Systems*, (Lake Tahoe, Nevada), pp. 548–556, 2012.
- [211] M. Richardson, R. Agrawal, and P. Domingos, *Trust management for the semantic web*, in *Proc. of International Semantic Web Conference (ISWC)*, pp. 351–368, Springer, 2003.
- [212] A. Olshevsky and J. N. Tsitsiklis, *On the nonexistence of quadratic Lyapunov functions for consensus algorithms*, *IEEE Trans. on Automatic Control* **53** (2008), no. 11 2642–2645.
- [213] M. B. Ye, J. Liu, B. D. Anderson, C. B. Yu, and T. Başar, *Evolution of social power in social networks with dynamic topology*, *IEEE Trans. on Automatic Control* (2018).
- [214] M. Kivela, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, *Multilayer networks*, *Journal of Complex Networks* **2** (2014), no. 3 203–271.
- [215] M. E. Newman and D. J. Watts, *Renormalization group analysis of the small-world network model*, *Physics Letters A* **263** (1999), no. 4-6 341–346.