Robust Peer-Monitoring on Graphs with an Application to Suicide Prevention in Social Networks*

Extended Abstract

Aida Rahmattalabi
University of Southern California
Los Angeles, CA
rahmatta@usc.edu

Phebe Vayanos
University of Southern California
Los Angeles, CA
phebe.vayanos@usc.edu

Anthony Fulginiti
University of Denver
Denver, CO
anthony.fulginiti@du.edu

Milind Tambe
University of Southern California
Los Angeles, CA
tambe@usc.edu

ABSTRACT

We consider the problem of selecting a subset of nodes (individuals) in a (social) network that can act as monitors capable of "watching-out" for their neighbors (friends) when the availability or performance of the chosen monitors is uncertain. Such problems arise for example in the context of "Gatekeeper Trainings" for suicide prevention. We formulate this problem as a two-stage robust optimization problem that aims to maximize the worst-case number of covered nodes. Our model is capable of incorporating domain specific constraints, e.g., fairness constraints. We propose a practically tractable approximation scheme and we provide empirical results that demonstrate the effectiveness of our approach.

KEYWORDS

Multiagent Systems; Social Network; Robust Optimization

ACM Reference Format:
Aida Rahmattalabi, Phebe Vayanos, Anthony Fulginiti, and Milind Tambe. 2019. Robust Peer-Monitoring on Graphs with an Application to Suicide Prevention in Social Networks. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

Motivation. Social networks are the substrate for a wide range of multiagent systems [1, 11, 12, 22]. In this regard, a notable problem consists in identifying, among all heterogeneous "actors" (or "agents") in the network, a suitable subset to most efficiently or effectively fulfill a particular task. In particular, we consider the problem of selecting a subset of nodes (individuals) in a (social) network that can act as "peer-monitors" capable of "watching-out" for their neighbors (e.g., friends) when the availability or performance of the chosen monitors is uncertain. Such problems arise for example in the context of "gatekeeper" trainings [15] which are considered best practice among suicide prevention professionals. Gatekeepers trained in suicide prevention and intervention are able to recognize warning signs of suicidal behavior among their peers and refer them to appropriate professionals as needed. Yet, when an individual is selected as a candidate to attend the training, they may be unwilling or unable to participate in the training or they may perform poorly as gatekeepers.

Traditional models and solutions for the monitoring problem on graphs under uncertainty assume that nodes are characterized solely by their position in the network, that the distribution of node availabilities is known, or only allow for budget constraints on the total number of monitors, see the literature review section.

These assumptions fail to hold in many real-world domains in which we are faced with heterogeneous agents (e.g., individuals with different races, genders, etc.). In such problems, and specially in socially sensitive domains such as gatekeeper training, it is crucial to ensure that agents of different types are treated fairly and equally (e.g., fairness by race). Finally, in many real-world contexts, very little information (historical data) is available to help inform or estimate the likelihood that a particular node will be available, should it be selected as a monitor. These characteristics motivate us to formulate the peer-monitoring problem on graphs under uncertainty in monitor performance as a robust optimization problem where we seek to maximize the worst-case coverage when some of the selected monitors do not succeed in the monitoring task (e.g., do not attend the training or do not perform well following training). Our contributions can be summarized as (i) modeling "robust peer-monitoring" problem as a novel two-stage robust optimization problem; (ii) providing a tractable approximation scheme, where the approximation level is adjustable by a single design parameter; (iii) a solution scheme that can handle arbitrary constraints of the domain, including fairness constraints.

Related Work. This problem is related to *robust sub-modular optimization.* In this regard, Orlin et al. [19] studied a a problem, in which one chooses a set of up to I items, and nature counteracts by eliminating at most J of those items. This work was followed by [7, 24], where in [24], the authors propose another greedy-based algorithm with a bound based on the curvature of the sub-modular function. Although these greedy algorithms are computationally efficient, they are not always implementable in practice due to domain specific constraints such as fairness restrictions.

*This article extends an earlier paper by [21].

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

Size (N)	Coverage (%)					Fairness Rate (%)				
	Exact	K=1	K=2	K=3	Greedy[24]	Exact	K=1	K=2	K=3	Greedy[24]
30	43.0	15.3	24.7	33.0	38.0	100	100	100	100	80
50	-	30.0	41.0	53.8	59.0	-	100	100	100	70
70	-	28.5	31.5	41.3	45.2	-	100	100	100	30

Table 1: Solution quality and fairness satisfiability of Exact, BD-forK (K=1,2,3) and a greedy heuristic, in real network samples

Our solution approach most closely relates to the robust optimization paradigm and specifically to the class of two-stage robust optimization with binary second-stage decisions. To approximate the solution to these problems, finite adaptability has been proposed. In particular, in [14], the authors show that for uncertainty sets as a bounded polyhedron, a two-stage robust optimization, can be approximately reformulated as an MILP. Later, [21] extended this result to a special case of discrete uncertainty sets. Finally, we note that robustness has been explored in different areas of multiagent systems as well. Notably, [10, 13, 18]. We note that this problem is conceptually different than the one we wish to investigate.

2 PROBLEM STATEMENT

A social network is represented as a directed graph $G = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ is the set of nodes (individuals) and $\mathcal{E}$ is the set of edges. An edge from u to v exists if v reports u as a friend. For a node u, $\delta(u) := \{v \in \mathcal{N} : (v, u) \in \mathcal{E}\}$ indicates its set of neighbors in G. We aim to maximize the coverage by selecting a set $\mathcal{X} \subseteq \mathcal{N}$ of at most I nodes. Next, "nature" selects a subset of at most $\lfloor \gamma I \rfloor$ nodes to be unavailable, where γ is the participation rate. A node is *covered* if at least one of its neighbors is chosen and available. Next, we provide the mathematical formulation of this problem.

2.1 Two-stage Robust Peer-Monitoring

First, we note that in the worst case, nature will choose the unavailable nodes from the set of chosen individuals, even if it is given the option to choose any node. This intuition allows us to formulate the peer monitoring problem as (1).

$$\max_{\boldsymbol{x} \in \mathcal{X}} \quad \min_{\xi \in \Xi} \max_{\boldsymbol{y} \in \{0,1\}^N} \left\{ \sum_{n \in \mathcal{N}} y_n : \sum_{n' \in \delta(n)} \xi_{n'} x_{n'} \geq y_n, \forall n \right\}, \quad (1)$$

in which Ξ is defined, independently of decision $\boldsymbol{x}$, as:

$$\Xi := \left\{ \xi \in \{0,1\}^{|\mathcal{N}|} \; : \; \sum_{n \in \mathcal{N}} (1 - \xi_n) \leq \lfloor \gamma I \rfloor \right\} \quad (2)$$

In this formulation, $\boldsymbol{x}$ is a binary vector and $x_n = 1$ if and only if node n is chosen. Also, $\xi_n = 1$ if and only if node n is available. Binary vector $\boldsymbol{y}$ indicates which nodes are covered. Set $\mathcal{X} = \{\boldsymbol{x} : \sum_{n \in \mathcal{N}} x_n \leq I\}$ is the set of all feasible node selections. We note this this set can also include arbitrary linear constraints. The first maximization problem decides which nodes $\boldsymbol{x}$ to invite. Following that, in the inner minimization problem, nature chooses the unavailable nodes. The set of all possible scenarios (nature's actions), is expressed as set Ξ. According to the definition of Ξ, at most $\lfloor \gamma I \rfloor$ can be unavailable. Finally, the inner-most maximization problem determines the covered nodes, where the constraints stipulate that a node can be covered if out of its chosen neighbors, at least one of them is available.

2.2 Solution Approach: K-Adaptability

K-Adaptability has been proposed to approximate the two-stage robust optimization problems with binary second-stage decisions. In K-Adaptability, K non-adjustable second-stage covering policies $\boldsymbol{y}^k, k \in \mathcal{K}$ ($\mathcal{K} := \{1, \cdots, K\}$) are chosen before the uncertain parameters are known. $y_n^k = 1$ indicates that node n is covered according to that policy. Once the node availabilities are observed, the best policy among the feasible ones will be implemented. Mathematically, the K-Adaptability counterpart problem is modeled as:

$$\max_{\substack{\boldsymbol{y}^k \in \{0,1\}^N \\ \boldsymbol{x} \in \mathcal{X}}} \min_{\xi \in \Xi} \max_{k \in \mathcal{K}} \Big\{ \sum_{n \in \mathcal{N}} y_n^k : y_n^k \leq \sum_{n' \in \delta(n)} \xi_{n'} x_{n'}, \forall n \Big\}. \quad (3)$$

In the K-adaptability problem, instead of considering each realization of the uncertain parameters, and finding the (robustly) optimal $\boldsymbol{y}$ for each scenario, we pre-commit to only K covering policies, decided in the first stage. Each policy encodes a possible network coverage, but it is decided before observing who among the chosen nodes is unavailable. The main contribution of this work is that we reformulate the K-adaptability problem exactly as a moderately sized mixed-integer linear program even though the set Ξ is discrete. We also, propose a Bender's decomposition based algorithm to efficiently solve this MILP formulation. Our algorithm augments the standard implementation with symmetry breaking cuts which significantly reduce the search of the K covering policies.

2.3 Results

We benchmark our approach against an exact approach, based on explicit enumeration of scenarios, and the *greedy* algorithm of [24], in terms of constraints satisfaction, and coverage. We evaluate on real social network samples of college students. We impose constraints to ensure different communities in the network will have equal representation in the gatekeeper training. We consider training capacity proportional to the network size ($B = 0.3N$), and failure rate equal to $\gamma = 0.3B$. The time limit is 2 hours for all approaches. Table 1 compares the number of times that the heuristics satisfy the fairness constraints. As expected, the K-adaptability solution is always feasible. In contrast, the heuristics are greatly challenged in the face of new domain constraints specially as the network size increases. We also investigate the performance of our algorithm for $K = 1, 2, 3$, where we observe we obtain competitive performance against greedy as the values of K increases. Finally, the exact approach does not scale beyond $N = 30$.

REFERENCES

[1] Aris Anagnostopoulos, Luca Becchetti, Carlos Castillo, Aristides Gionis, and Stefano Leonardi. 2012. Online Team Formation in Social Networks. In *Proceedings of the 21st International Conference on World Wide Web (WWW '12)*. ACM, New York, NY, USA, 839–848. https://doi.org/10.1145/2187836.2187950

[2] Rahmatollah Beheshti and Gita Sukthankar. 2014. A normative agent-based model for predicting smoking cessation trends. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, ACM, Paris, France, 557–564.

[3] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. 2009. *Robust optimization*. Vol. 28. Princeton University Press.

[4] Jacques F Benders. 1962. Partitioning procedures for solving mixed-variables programming problems. *Numerische mathematik* 4, 1 (1962), 238–252.

[5] Dimitris Bertsimas and Constantine Caramanis. 2010. Finite adaptability in multistage linear optimization. *IEEE Trans. Automat. Control* 55, 12 (2010), 2751–2766.

[6] Dimitris Bertsimas and Iain Dunning. 2016. Multistage robust mixed-integer optimization with adaptive partitions. *Operations Research* 64, 4 (2016), 980–998.

[7] Ilija Bogunovic, Slobodan Mitrović, Jonathan Scarlett, and Volkan Cevher. 2017. Robust submodular maximization: A non-uniform partitioning approach. *arXiv preprint arXiv:1706.04918* . (2017), –.

[8] Branislav Bošanskỳ, Albert Xin Jiang, Milind Tambe, and Christopher Kiekintveld. 2015. Combining compact representation and incremental generation in large games with sequential strategies. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. AAAI Press, MIT Press, Austin, Texas, 812–818.

[9] André Chassein, Marc Goerigk, Jannis Kurtz, and Michael Poss. 2018. Min-Max-Min Robustness for Combinatorial Problems with Budgeted Uncertainty. *arXiv preprint arXiv:1802.05072* . (2018), –.

[10] Chad Crawford, Zenefa Rahaman, and Sandip Sen. 2016. Evaluating the efficiency of robust team formation algorithms. In *International Conference on Autonomous Agents and Multiagent Systems*. Springer, ACM, Singapore, 14–29.

[11] Mathijs De Weerdt, Yingqian Zhang, and Tomas Klos. 2007. Distributed task allocation in social networks. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*. ACM, Honolulu, HI, USA, 76.

[12] Mathijs M de Weerdt, Yingqian Zhang, and Tomas Klos. 2012. Multiagent task allocation in social networks. *Autonomous Agents and Multi-Agent Systems* 25, 1 (2012), 46–86.

[13] Emir Demirovi, Nicolas Schwind, Tenda Okimoto, and Katsumi Inoue. 2018. Recoverable Team Formation: Building Teams Resilient to Change. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. ACM, Richland, SC, 1362–1370.

[14] Grani A Hanasusanto, Daniel Kuhn, and Wolfram Wiesemann. 2015. K-adaptability in two-stage robust binary programming. *Operations Research* 63, 4 (2015), 877–891.

[15] Michael Isaac, Brenda Elias, Laurence Y Katz, Shay-Lee Belik, Frank P Deane, Murray W Enns, Jitender Sareen, and Swampy Cree Suicide Prevention Team (12 members) 8. 2009. Gatekeeper training as a preventative intervention for suicide: a systematic review. *The Canadian Journal of Psychiatry* 54, 4 (2009), 260–268.

[16] Hiroaki Iwashita, Kotaro Ohori, Hirokazu Anai, and Atsushi Iwasaki. 2016. Simplifying urban network security games with cut-based graph contraction. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, ACM, Singapore, 205–213.

[17] Manish Jain, Dmytro Korzhyk, Ondřej Vaněk, Vincent Conitzer, Michal Pěchouček, and Milind Tambe. 2011. A double oracle algorithm for zero-sum security games on graphs. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, ACM, Taipei, Taiwan, 327–334.

[18] Tenda Okimoto, Nicolas Schwind, Maxime Clement, Tony Ribeiro, Katsumi Inoue, and Pierre Marquis. 2015. How to form a task-oriented robust team. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, ACM, Istanbul, Turkey, 395–403.

[19] James B Orlin, Andreas S Schulz, and Rajan Udwani. 2016. Robust monotone submodular function maximization. In *International Conference on Integer Programming and Combinatorial Optimization*. Springer, Waterloo, Canada, 312–324.

[20] Aida Rahmattalabi, Jen Jen Chung, Mitchell Colby, and Kagan Tumer. 2016. D++: Structural credit assignment in tightly coupled multiagent domains. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 4424–4429.

[21] Aida Rahmattalabi, Phebe Vayanos, and Milind Tambe. 2018. A Robust Optimization Approach to Designing Near-Optimal Strategies for Constant-Sum Monitoring Games. In *International Conference on Decision and Game Theory for Security*. Springer, Springer, Seattle, USA, 603–622.

[22] Liat Sless, Noam Hazon, Sarit Kraus, and Michael Wooldridge. 2014. Forming coalitions and facilitating relationships for completing tasks in social networks. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, ACM, Paris, France, 261–268.

[23] Anirudh Subramanyam, Chrysanthos E Gounaris, and Wolfram Wiesemann. 2017. K-Adaptability in Two-Stage Mixed-Integer Robust Optimization. *arXiv preprint arXiv:1706.07097* . (2017), .

[24] Vasileios Tzoumas, Konstantinos Gatsis, Ali Jadbabaie, and George J Pappas. 2017. Resilient monotone submodular function maximization. *arXiv preprint arXiv:1703.07280* . (2017), –.

[25] Phebe Vayanos, Daniel Kuhn, and Berç Rustem. 2011. Decision rules for information discovery in multi-stage stochastic programming. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*. IEEE, IEEE, Orlando, Florida, 7368–7373.

[26] Bryan Wilder, Han Ching Ou, Kayla de la Haye, and Milind Tambe. 2018. Optimizing network structure for preventative health. *International Conference on Autonomous Agents & Multiagent Systems* 15 (2018), 841–857.

[27] Bryan Wilder, Amulya Yadav, Nicole Immorlica, Eric Rice, and Milind Tambe. 2017. Uncharted but not Uninfluenced: Influence Maximization with an uncertain network. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, ACM, Sao Paulo, Brazil, 1305–1313.

[28] Amulya Yadav, Hau Chan, Albert Xin Jiang, Haifeng Xu, Eric Rice, and Milind Tambe. 2016. Using social networks to aid homeless shelters: Dynamic influence maximization under uncertainty. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, ACM, Singapore, 740–748.