Security Games with Contagion: Handling Asymmetric Information

(Extended Abstract)

Jason Tsai[1], Yundi Qian[1], Yevgeniy Vorobeychik[2], Christopher Kiekintveld[3], Milind Tambe[1]

[1]University of Southern California, Los Angeles, CA 90089
{jasontts, yundi.qian, tambe} @usc.edu
[2]Sandia National Laboratories, Livermore, CA
yvorobe@sandia.gov
[3]University of Texas at El Paso, El Paso, TX
cdkiekintveld@utep.edu

ABSTRACT

Counterinsurgency, which is the effort to mitigate support for an opposing organization, is one such domain that has been studied recently and past work has modeled the problem as an influence blocking maximization that features an influencer and a mitigator. While past work has introduced scalable heuristic techniques for generating effective strategies using a double oracle algorithm, it has not addressed the issue of uncertainty and asymmetric information, which is the topic of this paper.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence

General Terms

Algorithms, Security, Performance

Keywords

Game theory, Social contagion, Influence maximization

1. INTRODUCTION

The spread of information and social behaviors has been studied extensively in many disciplines in the context of phenomena such as viral marketing, rumor spreading and the Arab Spring [9, 12, 13]. Counterinsurgency, the competition for the support of local leadership, has also been studied as a game with two strategic players [6, 5, 14]. Although many aspects of this problem are highly active areas of research, the key *computational* question is to decide which local leaders to influence to achieve each player's primary goal: maximize influence for one player, and mitigate the first player's influence for the other player.

These 'counter-contagion' games have received recent attention in the security games literature [14] and has been modeled as a graph where nodes represent leaders and edges between the nodes representing the probability of influence. This line of research, however, has not yet examined the impact of asymmetric information. Informational challenges abound in counterinsurgency, where the insurgents are typically an indigenous group that has an informational advantage and the mitigators often have uncertainty about their knowledge of the social network [6]. Figure 1, for example, shows a realistic social network for the leadership of a set of local villages in Afghanistan [6]. Given real-world information constraints, the counterinsurgency team may not have perfect information of the graph and be uncertain about some set of edges.

In our work, the mitigator's uncertainty about the graph structure is modeled as a Bayesian game with each Bayesian type representing a separate instantiation of the graph. The mitigator's strategy must now reason over the distribution of types. The influencer's (insurgent's) perfect knowledge of the graph structure allows him to specify a behaviorial strategy which conditions the strategy used on the specific type. This results in Bayesian games with an exponential number of types where each game is already extremely challenging to solve efficiently. We aim to address this class of problems efficiently and effectively.

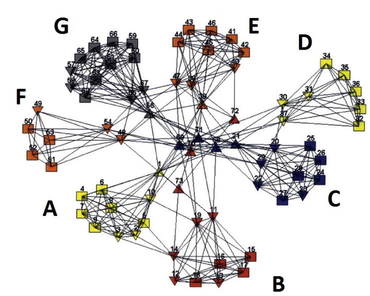

Figure 1: Example Afghani leadership network

2. RELATED WORK

Recent work in game-theoretic security allocation have also dealt with domains that were modeled as graphs [1, 7, 4], however their actions were all deterministically defined and did not feature a probabilistic contagion component. The work in uncertainty in security games is also relevant [8, 10, 15], but once again do not feature the contagion component found in our domain.

This contagion process has been studied outside of the security

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

games literature and is known as influence maximization, in which a player attempts to optimize a selection of beginning 'seed' nodes from which to spread his influence in a known graph. This class of problems were first introduced as a discrete maximization problem by Kempe et al. (2003) who showed submodularity of the maximization problem, enabling a greedy approximation. This work has been followed-up by numerous proposed speed-up techniques [3, 11, 12].

Two-player variants of influence maximization have been studied as well, one of which is known as influence blocking maximization problems and are equivalent to the counter-contagion games we study. These models have been explored with both independent cascade and linear threshold models of propagation [2, **?**], however, work in this area has generally focused only on the defender's best-response problem. The exception is Tsai et al. [14] which addresses the algorithmic challenge of finding equilibria strategies. Finally, Hung et al. (2011) and Howard (2010) also model counterinsurgency and attempt to optimize against a strategic adversary. However, none of these works model the uncertainty that is critical in domains such as counterinsurgency.

3. ASYMMETRIC INFORMATION GAME

We model counterinsurgency as a two-player Bayesian zero-sum game situated on a graph in which two players, the influencer (denoted by I) and the mitigator (denoted by M) compete to maximize influence over the nodes. Formally, let $G = (V, E)$ be a graph with weighted nodes V and edges E, and for each edge $(i, j) \in E$, let p_{ij} be the probability that node i's opinion will influence node j. We model propagation of influence in the graph as a synchronized independent cascade process [9] as follows. Suppose that the influencer initially attempts to influence a subset of nodes $S_I \subseteq V$ to his cause, and the mitigator's initial influence is aimed at a subset of nodes $S_M \subseteq V$. For nodes $v \in S_I \cap S_M$ which both players initially try to influence, initial "activation" (e.g., actual opinion adoption) happens in either player's favor with equal probability, while all the remaining nodes adopt the view of (are activated by) the player who directly targets them. Next, we activate all edges (i, j) in the graph with the corresponding probability of influence, p_{ij}. At that point, the influence process proceeds through a sequence of iterations. In each iteration, if a node j has not yet adopted an opinion but has active edges to neighbors who have, j either adopts the opinion of these neighbors when it is unanimous, or adopts each opinion with equal probability if j's active neighbors disagree. Viewing now the initial target nodes S_I and S_M as the strategies of the players I and M respectively, let $\sigma(S_I, S_M)$ be the expected value of nodes that adopt the influencer's opinion following the independent cascade process described above. We define the utility of the influencer to be $U_I(S_I, S_M) = \sigma(S_I, S_M)$.

Our model differs from those in past works (e.g., Tsai et al. 2012) by relaxing the complete/symmetric information assumption. Specifically, we assume that the influencer knows the actual influence graph G exactly, while the mitigator is uncertain about its true structure, and only knows the probability distribution over possible graphs. Let λ be an index identifying a particular graph G_λ, and let us make explicit the dependence of the expected influence on the graph, denoting it by $\sigma(S_I, S_M, \lambda)$. Finally, we denote by P the probability distribution over λ, with P_λ the probability that the true graph is the one identified by λ. From the mitigator's perspective, the influencer's decision will depend on his type, that is, on the true graph which the influencer observes. Thus, we view the influencer's strategy S_I as a function of λ, with S_I^λ denoting the influencer's strategy when his type is λ. The mitigator's utility is then $U_M(S_I, S_M) = -E_{\lambda\sim P}[\sigma(S_I^\lambda, S_M, \lambda)]$.

4. CHALLENGE

Though these Bayesian counter-contagion games are zero-sum and, therefore, amenable to linear programming solutions, the asymmetric information component adds a major dimension of difficulty. Specifically, because the uncertainty occurs over graph instances, the number of influencer types can be exponentially large. Since each individual game is already very challenging to solve as per Tsai et al. 2012, exponentially many of them exacerbates this challenge. Handling this uncertainty efficiently and effectively remains a major open challenge to real-world application of these techniques.

5. REFERENCES

[1] N. Basilico and N. Gatti. Automated abstractions for patrolling security games. In *AAAI*, 2011.

[2] C. Budak, D. Agrawal, and A. E. Abbadi. Limiting the spread of misinformation in social networks. In *WWW*, pages 665–674, 2011.

[3] W. Chen, C. Wang, and Y. Wang. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *KDD*, pages 1029–1038, 2010.

[4] E. Halvorson, V. Conitzer, and R. Parr. Multi-step multi-sensor hider-seeker games. In *IJCAI*, pages 159–166, 2009.

[5] N. J. Howard. *Finding optimal strategies for influencing social networks in two player games*. Masters thesis, MIT, Sloan School of Management, June 2011.

[6] B. W. K. Hung. *Optimization-Based Selection of Influential Agents in a Rural Afghan Social Network*. Masters thesis, MIT, Sloan School of Management, June 2010.

[7] M. Jain, D. Korzhyk, O. Vanek, V. Conitzer, M. Pechoucek, and M. Tambe. A double oracle algorithm for zero-sum security games on graphs. In *AAMAS*, pages 327–334, 2011.

[8] M. Jain, M. Tambe, and C. Kiekintveld. Quality-bounded solutions for finite bayesian stackelberg games: Scaling up. In *International Conference on Autonomous Agents and Multiagent Systems*, 2011.

[9] D. Kempe, J. M. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *KDD*, pages 137–146, 2003.

[10] C. Kiekintveld, J. Marecki, and M. Tambe. Approximation methods for infinite bayesian stackelberg games: Modeling distributional payoff uncertainty. In *International Conference on Autonomous Agents and Multiagent Systems*, 2011.

[11] M. Kimura, K. Saito, R. Nakano, and H. Motoda. Extracting influential nodes on a social network for information diffusion. *Data Min. Knowl. Discov.*, 20(1):70–97, 2010.

[12] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. M. VanBriesen, and N. S. Glance. Cost-effective outbreak detection in networks. In *KDD*, pages 420–429, 2007.

[13] M. Trusov, R. E. Bucklin, and K. Pauwels. Effects of word-of-mouth versus traditional marketing: Findings from an internet social networking site. *Journal of Marketing*, 73, September 2009.

[14] J. Tsai, T. H. Nguyen, and M. Tambe. Security games for controlling contagion. In *AAAI*, 2012.

[15] Z. Yin and M. Tambe. A unified method for handling discrete and continuous uncertainty in bayesian stackelberg games. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.