Mitigating the Curse of Correlation in Security Games by Entropy Maximization

Extended Abstract

Haifeng Xu, Shaddin Dughmi, Milind Tambe, Venil Loyd Noronha
University of Southern California
Los Angeles, CA
{haifengx,shaddin,tambe,vnoronha}@usc.edu

ABSTRACT

In Stackelberg security games, a defender seeks to randomly allocate limited security resources to protect critical targets from an attack. In this paper, we study a fundamental, yet underexplored, phenomenon in security games, which we term the *Curse of Correlation* (CoC). Specifically, we observe that there are *inevitable* correlations among the protection status of different targets. Such correlation is a crucial concern, especially in *spatio-temporal* domains like conservation area patrolling, where attackers can surveil patrollers at certain areas and then infer their patrolling routes using such correlations. To mitigate this issue, we propose to design entropy-maximizing defending strategies for spatio-temporal security games, which frequently suffer from CoC. We prove that the problem is #P-hard in general. However, it admits efficient algorithms in well-motivated special settings.

KEYWORDS

Security Games; Correlation; Information Leakage; Max Entropy

ACM Reference Format:
Haifeng Xu, Shaddin Dughmi, Milind Tambe, Venil Loyd Noronha. 2018. Mitigating the Curse of Correlation in Security Games by Entropy Maximization. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

The security game is played between a defender and an attacker where the defender's goal is to *randomly* allocate a limited number of security resources to protect targets from attack [6, 11, 14]. Standard models assume that the attacker only knows the defender's randomized strategy, but is oblivious to the sampled real-time allocation. However, this assumption may fail since in many situations the attacker can *partially* observe the defender's real-time allocation. Such partial observation can be utilized to infer extra information about the overall strategy. This is particularly the case in games played out in space and time, a.k.a. *spatio-temporal* security games, which are also the primary focus of this work. For example, it has been reported that in wildlife protection domains, some poachers partially monitor rangers' patrolling activities and then make their poaching plans based on their observations [9, 10]. Similar issues could happen when optimizing the patrolling on graphs [2, 5] and scheduling of air marshals [1].

We observe that the randomized allocation of limited security resources creates *inherent* correlations within protection statuses of different targets – the coverage of some targets implies that other targets are not protected. Such correlation allows the attacker to utilize his partial observation at some targets to infer information about other targets' protection status, a phenomenon which we term the *Curse of Correlation* (CoC). We show that such correlation is inevitable, and may cause significant loss if not addressed properly. To overcome this challenge, we propose to adopt the "most random" defending strategy, or more formally, the strategy that, subject to optimizing the usual objective under required constraints, maximizes (Shannon) *entropy*. Intuitively, such a strategy could be resistant to partial leakage due to its extreme randomness/unpredictability.

To that end, this paper offers the following contributions. First, we formally study the Curse of Correlation (CoC) phenomenon in security games and illustrate the importance of handling CoC, particularly in spatio-temporal domains. Second, we propose to adopt the defending strategy with maximum entropy and illustrate its advantages when compared to the idealized optimal solution tailored to a specific leakage model. Third, we design novel algorithms to sample defending strategies with maximum entropy in spatio-temporal security games. We prove that the exact max-entropy defending strategy is #P-hard to compute in general, but admits polynomial-time algorithms in well-motivated special settings as well as efficient heuristic algorithms.

2 MOTIVATING EXAMPLE

Our example is about the protection of conservation areas, though the phenomenon it illustrates could occur in any security setting. The problem concerns designing rangers' patrolling schedules within a fixed time period, say, a day. This is usually modeled by discretizing the area into cells as well as discretizing the time. At the top of Figure 1, we depict a concrete example with 4 cells to be protected at 3 time layers: morning, noon and afternoon. The numbers around each cell are the desired marginal coverage probabilities for each cell at each time (color thickness depicts the probability density). The defender has 2 rangers, and seeks to randomize their patrolling to achieve the required marginal probabilities.

To deploy a mixed strategy of small support, as done by traditional algorithms, one can implement the marginal vector by mixing the three pure strategies listed at the bottom of Figure 1 (filled dots are fully covered). Unfortunately, it turns out that such

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

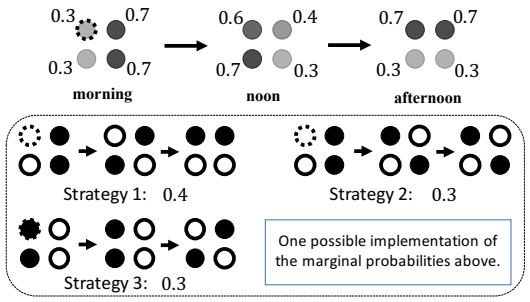

Figure 1: Marginal and Implementation.

an implementation is extremely vulnerable to the attacker's partial surveillance. For example, if the attacker can surveil the status of the top-left cell in the morning (i.e., the one with dashed boundary) and prepare an attack in the afternoon, he can always find a completely uncovered cell to attack. Specifically, if the dashed cell is covered, this means `Strategy` 3 is deployed and two cells will be uncovered in the afternoon; Otherwise, either `Strategy` 1 or 2 is deployed, and the bottom left cell will be uncovered in the afternoon. So the attacker can successfully identify uncovered cells in the afternoon by monitoring only *one* cell in the morning.

The issue above is due to the inherent correlation among the protection status of different targets when allocating a limited number of resources. Particularly, the coverage of some targets must imply that some other targets are unprotected. The example illustrates how the attacker can take advantage of such correlation and infer a significant amount of information about the protection of other targets by monitoring even a single target. This is what we term the *Curse of Correlation* (CoC) in security games.

3 REMEDY BY ENTROPY MAXIMIZATION

To tackle the curse of correlation in security games, the ideal approach is to come up with an accurate model to capture the attacker's partial observation (also referred to as an information *leakage model* henceforth for convenience), and then solve the model to obtain the defender's optimal defending strategy. However, this approach suffers from several critical drawbacks. First, it is usually very difficult to obtain an accurate leakage model in reality since the attacker's choice of target monitoring depends on many hidden factors, thus is highly unpredictable. Second, even if the defender has an accurate leakage model, computing the optimal defender strategy against the leakage model is usually intractable [13]. Third, another concern about any optimal solution tailored to a specific leakage model is that such a solution may be easily "gamed" by the attacker. In particular, the optimal solution naturally biases towards the leaking targets by assigning more security forces to these targets. This however opens the door for the attacker to strategically manipulate the defender's belief on leaking targets, e.g., by intentionally spreading misleading information, with the goal of shifting the defender's defense away from the attacker's prime targets.

Entropy maximization – a more robust solution

These barriers motivate our adoption of a more robust (though inevitably more conservative) approach. Particularly, we propose to first compute the optimal defender strategy assuming no leakage and then adopt the max-entropy implementation of its marginal vector. Our choice of max entropy is due to at least three reasons. First, the max-entropy strategy is the most random, thus unpredictable, defender strategy. We believe that this is a natural choice when the defender is uncertain about which target leaks information (the setting we are in). Second, the max-entropy distribution exhibits substantial approximate stochastic independence among the protection statuses of targets[1], so that the protection status of any leaking targets does not carry much information about that of others. Third, experiments show that our max entropy approach performs extremely well in comparisons with several other alternatives; in fact, in some settings, it achieves a solution quality that is even close to the optimal defender utility under no leakage! Given such significant empirical results, entropy maximization clearly stood out as a powerful approach to address information leakage.

Moreover, the max-entropy approach also enjoys several practical advantages. First, it does not require a concrete leakage model. Instead, it seeks to reduce the overall correlation among the statuses of all targets, thus serves as a robust solution. Second, this approach is "compatible" with currently deployed security systems which assume no leakage, since it preserves their marginal coverage probabilities and only "adds" additional randomness to protect against leakage; in this sense, our approach strictly improves the solution. This is particularly useful in domains where building a new security system is not feasible or too costly.

Computing the Max-Entropy Distribution

Efficient computation remains a challenge for our approach due to the widely-known difficulty of computing the max-entropy distribution over combinatorial structures subject to given marginal probabilities. Indeed, we show that the problem is #P-hard in general for spatio-temporal security games. Fortunately, we are able to prove that the max-entropy distribution can be computed in polynomial time (in the input size of the problem) for two well-motivated security settings: (1) constant number of security resources; (2) the air marshal scheduling problem with round trips. Our algorithms employ sampling techniques, duality theory and also rely crucially on the structure of the problem.

4 RELATED WORK.

Alon et al. [3] study information leakage in *normal-form* zero-sum games and exhibit NP-hardness results in several model variants. This work also relates to the line of research on adversarial patrolling games (APGs) [2, 4, 5, 7, 12]. APGs also consider the attacker's real-time observations, however our settings differ from APGs in several aspects: 1. the defender in APGs *typically* has one patroller and the attacker has *full knowledge* of the defender's movements, while in our setting the defender has many security resources and the attacker only observes a small subset of targets; 2. APGs assume that the attacker takes time to complete an attack, while attacks in our setting are instantaneous. These important differences make the algorithms for APGs inapplicable to our settings.

Acknowledgment: This research was supported by MURI grant W911NF-17-1-0370, NSF grant CCF-1350900 and a Google PhD Fellowship.

[1]This is widely observed in practice, and also theoretically proved in some settings, e.g., matchings [8].

REFERENCES

[1] ABC News. May 22, 2006. Congress: Terrorists Could Spot Undercover Air Marshalsa. (May 22, 2006).

[2] Noa Agmon, Vladimir Sadov, Gal A. Kaminka, and Sarit Kraus. 2008. The Impact of Adversarial Knowledge on Adversarial Planning in Perimeter Patrol, Vol. 1. 55–62.

[3] Noga Alon, Yuval Emek, Michal Feldman, and Moshe Tennenholtz. 2013. Adversarial leakage in games. *SIAM Journal on Discrete Mathematics* 27, 1 (2013), 363–385.

[4] Steve Alpern, Alec Morton, and Katerina Papadaki. 2011. Patrolling games. *Operations research* 59, 5 (2011), 1246–1257.

[5] Nicola Basilico, Nicola Gatti, and Francesco Amigoni. 2012. Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder. *Artificial Intelligence* 184 (2012), 78–123.

[6] Avrim Blum, Nika Haghtalab, and Ariel D Procaccia. 2014. Learning optimal commitment to overcome insecurity. In *NIPS*.

[7] Branislav Bošanský, Viliam Lisý, Michal Jakob, and Michal Pěchouček. 2011. Computing time-dependent policies for patrolling games with mobile targets.. In *AAMAS*. IFAAMAS.

[8] Jeff Kahn and P Mark Kayll. 1997. On the stochastic independence properties of hard-core distributions. *Combinatorica* 17, 3 (1997), 369–391.

[9] William Moreto. 2013. *To conserve and protect: Examining law enforcement ranger culture and operations in Queen Elizabeth National Park, Uganda.* Ph.D. Dissertation. Rutgers University-Graduate School-Newark.

[10] Vincent R Nyirenda and Chansa Chomba. 2012. Field foot patrol effectiveness in Kafue National Park, Zambia. *Journal of Ecology and the Natural Environment* 4, 6 (2012), 163–172.

[11] Milind Tambe. 2011. *Security and game theory: algorithms, deployed systems, lessons learned.* Cambridge University Press.

[12] Yevgeniy Vorobeychik, Bo An, and Milind Tambe. 2012. Adversarial patrolling games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*. International Foundation for Autonomous Agents and Multiagent Systems, 1307–1308.

[13] Haifeng Xu, Albert X. Jiang, Arunesh Sinha, Zinovi Rabinovich, Shaddin Dughmi, and Milind Tambe. 2015. Security Games with Information Leakage: Modeling and Computation. In *IJCAI*.

[14] Yue Yin, Haifeng Xu, Jiarui Gain, Bo An, and Albert Xin Jiang. 2015. Computing optimal mixed strategies for security games with dynamic payoffs. In *IJCAI*.