Adaptive and Dynamic Allocation of Security Resources (Doctoral Consortium)

Sara Marie Mc Carthy
University of Southern California
sara.m.mccarthy@gmail.com

ABSTRACT

This research aims to address the challenge of dynamic resource allocation, where security resources are heterogeneous and the defender is tasked with simultaneously optimizing over their investment into these resources in order to build the best possible team, as well as their deployment in the field. This allows the defender to design more adaptive strategies, tailoring the set of resources used to the specific parameters of a domain, and allowing them to adapt to the realization of uncertain or time varying parameters.

Keywords

Game Theory, Team Formation

1. CURRENT RESULTS

1.1 Team Formation

Past work in security games has focused on the problem of static resource allocation; how to optimally deploy a given fixed team of resources. While the common challenge here is the optimization of a limited set of resources, it is often the case that budget for investment in these resources is also limited. The challenge arises when we ask the question of how to optimize this investment. Given a set of resources, each with varying costs, and effectiveness, with the ability to coordinate in the field and we want to be able to build the best possible team of these resources and compute the optimal deployment of this team. This optimization over a portfolio of security resources is particularly challenging given that the expected return on such an investment depends on the actual deployment of these resources in the field.

This work provides a formal model of this problem, which we refer to as the Simultaneous Optimization of Resource Teams and Tactics (SORT) as a new fundamental research problem in security games that combines strategic and tactical decision making [3]. I studied this problem using the challenge of optimizing the defense of forests against illegal logging in Madagascar as a motivating domain, where I combine two different classes of security games, Green Security and Network Security Games, where the challenge of simultaneously optimizing teams and deployment is particularly challenging. Green Security Games [4] [7] concern themselves with the protection of forests, fish and wildlife, a critical problem in environmental sustainability. They provide a unique domain space where we often we are faced with the challenge of teaming up many different groups, from national police to forest guards to NGOs, each with differing capabilities and costs, making the strategic optimization challenging. Unfortunately these often occur in developing countries, and budgets for protecting these natural resources are often very limited, making it crucial to allocate security resources efficiently. Network Security Games provide a challenge for the tactical decision making problem as there is combinatorially large space of defender strategies which need to be optimized over [6]. The fact that the tactical question is computationally challenging emphasizes the difficulty of the SORT problem, which requires evaluating the effectiveness of many teams to select the right one.The challenge of optimizing a team of security resources can be formulated as the following optimization problem, $\max_{\lambda \subset R} \left\{ F(\lambda) : \sum_{k \in \lambda} b_k \leq B \right\}$ where the value of a team of resources λ selected from some set of resources R is given by the expected utility of their optimal deployment, denoted $F(\lambda)$.

$F(\lambda)$ can computationally difficult to calculate, particularly in Network Security Games, because it requires finding the optimal tactical allocation to assess the utility of a given team λ. Since there are an exponentially many possible teams, the sequential approach of evaluating $F(\lambda)$ exactly for every team and picking the best one is impractical. Exactly evaluating $F(\lambda)$ requires solving the underlying network security game, which can formulated as an optimization problem and is solved using a double oracle algorithm. Instead, in our approach to SORT, I developed FORTIFY (Forming Optimal Response Teams for Forest safetY), a scalable branch and bound style algorithm which integrates the analysis of the strategic and tactical aspects of the problem to efficiently approximate $F(\lambda)$ and search the space of teams much more efficiently and limit the number of instances where we evaluate $F(\lambda)$ exactly. FORTIFY uses a hierarchical approach to abstract the Network Security Game at varying levels of detail, providing bounds on $F(\lambda)$ the value of different teams of resources to speed up the search for the optimal team. The novelty is how we generate bounds on different teams, using successive relaxations of the security game problem, which is the key component to the algorithm. The relaxations are made in such a way that they will always provide upper bounds on the true team value and will be easier to solve than the full security game. We iteratively refine these relaxations, allowing us to generate tighter and tighter upper bounds, until a team is finally evaluated using the full game model. I evaluated this work, using real network and resource data obtained from our partners working in Madagascar, and show that not only is the algorithm scalable, but as the number of possible combination of resources grows it because more and more worthwhile to perform this optimization over resources.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

1.2 Active Sensing

I also studied the problem of active sensing in a computer network, where the defender observes a stream of suspicious alerts (or absence of alerts), and is tasked with inferring if an attack is taking place, and determining the best response policy. While the defender may have many kinds of security resources employed to monitor network traffic, such resources are inherently noisy and each single alert does not provide a high confidence estimate about the security state [8][5]. Thus, the defender needs to come up with a sequential plan of actions while dealing with uncertainty in the network and in the alerts, and must weigh the cost of deploying detectors to increase their knowledge with the potential loss due to successful attacks as well as the cost of misclassifying legitimate network use. This problem of active sensing is common to a number of cyber security problems; here I focus on the challenge of detecting and addressing advanced persistent threats (APTs), with particular focus on data exfiltration over DNS queries. I developed a decision-theoretic planning model to reason about noisy observations in order to dynamically allocate security resources such as sensors and determine whether or not the suspicious activity is malicious, and compute the best response policy.

More specifically, I provide a novel Virtually Distributed Partially Observable Markov Decision Process (VD-POMDP) formulation to address the challenge [2], where the efficiency of the formulation is based on two key contributions: (i) the problem is decomposed in a way that allows for individual sub-POMDPs with sparse interactions to be created. Individual policies for different sub-POMDPs are planned separately and their sparse interactions are only resolved at execution time to determine the joint actions to perform; (ii) The abstraction in planning step allows for large speedups and scalability, after which a fast MILP is used to implement the abstraction while resolving any interactions. This allows us to determine optimal sensing strategies, leveraging information from many noisy detectors, and subject to constraints imposed by network topology, forwarding rules and performance costs on the frequency, scope and efficiency of sensing we can perform. I provide conditions under which our methods are guaranteed to result in the optimal joint policy, and provide empirical evidence to show that the final policy still performs well when these conditions are not satisfied. I also provide experimental evaluation of our model in a real network testbed, where I demonstrate the ability to correctly identify real attacks.

2. RESEARCH PLAN

2.1 Adaptive Screening

I am currently working on the challenge of adaptive screening in Threat Screening Games (TSG). Previous work in TSGs [1] operates under the strong assumption of perfect knowledge of the future world states; in particular it assumes that the exact number of screenees arriving in each time window is known a priori. However, this assumption does not hold as the distribution of passenger arrivals at the screening checkpoint by time window and screenee category is inherently uncertain. When these uncertain parameters are realized in a way that is different that what had been planned for, this can result in not only sub-optimal screening strategies, but can also result in strategies that are not implementable. I am working on relaxing this assumption, by introducing a new model of Uncertain Threat Screening Games (UTSG) which allows us to explicitly model the uncertainty, and develop robust screening strategies in the face of such uncertainty, where we plan for any possible realization of screenee arrival times. This has an important consequence on the model given the sequential nature of the game, as we can no longer reason about each time window individually as with the previous model of TSG. The screenee arrival information observed as well as our decision in previous time windows can be used to reason about the uncertainty in the remaining time windows and thus modify the screening strategy going forward.

The defenders optimal strategy is then a policy, an adaptive screening strategy determined not only by their interactions with the adversary, but also conditioned on the realization of the uncertain parameters of the problem; the arrival of passengers in previous time windows. The defender then needs to come up with a valid strategy for every possible realization of the uncertainty in arrivals, so that a major challenge in this work is in dealing with the huge space of possible passenger arrivals.

2.2 Future Work

Looking forward, I am currently interested in exploring problem domains where AI may be used to address challenges in social welfare and social good. I plan to expand on my work in the domain of green security and look at improving the scalability of my solution methods to the SORT problem, addressing the challenging case where it is infeasible to consider the entire space of team which may be formed. I am also looking at how learning may be used for optimization, in particular at how deep neural networks can be used to help solve the large scale optimization problems that arise from adaptive planning strategic decision making with a dynamic set of resources in large scale games.

REFERENCES

[1] M. Brown, A. Sinha, A. Schlenker, and M. Tambe. One size does not fit all: A game-theoretic approach for dynamically and effectively screening for threats. In *AAAI conference on Artificial Intelligence (AAAI)*, 2016.

[2] S. M. Carthy, A. Sinha, M. Tambe, and P. Manadhata. Data exfiltration detection and prevention: Virtually distributed pomdps for practically safer networks. In *Decision and Game Theory for Security (GameSec 2016)*, 2016.

[3] S. M. Carthy, M. Tambe, C. Kiekintveld, M. L. Gore, and A. Killion. Preventing illegal logging: Simultaneous optimization of resource teams and tactics for security. In *AAAI conference on Artificial Intelligence (AAAI)*, 2016.

[4] F. Fang, P. Stone, and M. Tambe. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.

[5] G. FarnHam. Detecting dns tunneling. Technical report, SANS Institute InfoSec Reading Room, Februrary 2013.

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

[7] M. P. Johnson, F. Fang, , and M. Tambe. Patrol strategies to maximize pristine forest area. In *Conference on Artificial Intelligence (AAAI)*, 2012.

[8] R. Sommer and V. Paxson. Outside the closed world: On using machine learning for network intrusion detection. In *Security and Privacy (SP), 2010 IEEE Symposium on*, pages 305–316. IEEE, 2010.